


OXFORD

A complex network visualization composed of numerous thin, glowing lines of various colors (orange, red, blue, green) against a black background. The lines form a dense web of connections, with some nodes appearing as bright, glowing points of light. The overall effect is organic and suggests a complex system of interactions.

THE CONNECTED PAST

Challenges to Network Studies
in Archaeology and History

Edited by Tom Brughmans, Anna Collar, and Fiona Coward

THE CONNECTED PAST

The Connected Past

*Challenges to Network Studies
in Archaeology and History*

Edited by
TOM BRUGHMANS, ANNA COLLAR,
AND FIONA COWARD

OXFORD
UNIVERSITY PRESS



Great Clarendon Street, Oxford, OX2 6DP,
United Kingdom

Oxford University Press is a department of the University of Oxford.
It furthers the University's objective of excellence in research, scholarship,
and education by publishing worldwide. Oxford is a registered trade mark of
Oxford University Press in the UK and in certain other countries

© Oxford University Press 2016

The moral rights of the authors have been asserted

First Edition published in 2016

Impression: 1

All rights reserved. No part of this publication may be reproduced, stored in
a retrieval system, or transmitted, in any form or by any means, without the
prior permission in writing of Oxford University Press, or as expressly permitted
by law, by licence or under terms agreed with the appropriate reprographics
rights organization. Enquiries concerning reproduction outside the scope of the
above should be sent to the Rights Department, Oxford University Press, at the
address above

You must not circulate this work in any other form
and you must impose this same condition on any acquirer

Published in the United States of America by Oxford University Press
198 Madison Avenue, New York, NY 10016, United States of America

British Library Cataloguing in Publication Data
Data available

Library of Congress Control Number: 2015943401

ISBN 978-0-19-874851-9

Printed in Great Britain by
Clays Ltd, St Ives plc

Links to third party websites are provided by Oxford in good faith and
for information only. Oxford disclaims any responsibility for the materials
contained in any third party website referenced in this work.

Contents

<i>List of Figures</i>	vii
<i>List of Tables</i>	xi
<i>List of Contributors</i>	xiii

PART I: CHALLENGING NETWORK METHODS AND THEORIES

1. Network Perspectives on the Past: Tackling the Challenges <i>Tom Brughmans, Anna Collar, and Fiona Coward</i>	3
2. Networks in Archaeology: Between Scientific Method and Humanistic Metaphor <i>Carl Knappett</i>	21
3. Networks or Work-nets? Actor-Network Theory and Multiple Social Topologies in the Production of Roman <i>Terra Sigillata</i> <i>Astrid Van Oyen</i>	35

PART II: CHALLENGING NETWORK ANALYSIS OF ARCHAEOLOGICAL AND HISTORICAL DATA

4. Analytical Challenges for the Application of Social Network Analysis in Archaeology <i>Matthew A. Peeples, Barbara J. Mills, W. Randall Haas, Jr, Jeffery J. Clark, and John M. Roberts, Jr</i>	59
5. How Reliable are Centrality Measures for Data Collected from Fragmentary and Heterogeneous Historical Sources? A Case Study <i>Marten Düring</i>	85
6. Uncovering the Hidden Routes: Algorithms for Identifying Paths and Missing Links in Trade Networks <i>Constantinos Tsirogiannis and Christos Tsirogiannis</i>	103

PART III: CHALLENGING NETWORK MODELS

7. Can Archaeological Models Always Fulfil our Prejudices? <i>Ray Rivers</i>	123
---	-----

8. Which Network Model Should I Use? Towards a Quantitative Comparison of Spatial Network Models in Archaeology <i>Tim Evans</i>	149
9. Networks, Homophily, and the Spread of Innovations <i>Anne Kandler and Fabio Caccioli</i>	175
<i>Index</i>	199

List of Figures

1.1. An abstract representation of a network model and an example of how a network model is used to explore a particular phenomenon.	8
1.2. Different data types representing individuals and social relationships.	9
4.1. Map of the Southwest Social Networks Project study area.	61
4.2. Overview of the ceramic apportioning procedure.	68
4.3. Network diagrams for the southern Southwest for the AD 1300–1350 and AD 1350–1400 intervals.	70
4.4. Network diagram of the greater Hohokam region for the AD 1250–1300, 1300–1350, and 1350–1400 intervals.	72
4.5. Network diagrams of the entire Southwest Social Networks Project study area for the AD 1250–1300, 1300–1350, 1350–1400, and 1400–1450 intervals.	73
4.6. Correlations between degree centrality in full network for each 50-year interval and sub-samples at sampling fractions from 90% to 10% of total.	76
4.7. Correlations between eigenvector centrality in full network for each 50-year interval and sub-samples at sampling fractions from 90% to 10% of total.	76
5.1. Matches between influential actors and centrality scores.	96
5.2. Mapping influential actors within their networks.	98
6.1. The undirected version of the network extracted from Watson and Todeschini (2007).	107
6.2. The directed version of the network extracted from Watson and Todeschini (2007).	108
6.3. The path subnetwork between the node that corresponds to Robert Hecht and the node that corresponds to the Getty Museum.	110
6.4. The performance of the two methods that indicate missing links in the network.	118
7.1. Important MBA sites for the Southern Aegean.	136
7.2. Testing the sensitivity of the simple gravity model and the doubly constrained gravity model to the distance scale for single journeys.	138

7.3. Exemplary <i>ariadne</i> networks for $D = 100$ km before and after the eruption of Thera.	140
7.4. The 109 archaic settlements based on figure 1 from Rihll and Wilson (1987) and using the same labels.	142
8.1. Four Minoan sites used for examples, indicating the direct routes to and from Knossos and the quickest routes.	151
8.2. Two simple networks derived using the quickest distances: the MDN network for a 110 km cutoff, and the PPA $k=1$ network.	154
8.3. Illustration of the flow F_{st} from site s to site t .	157
8.4. Illustration of typical forms for the deterrence function f_{st} as a function of the distance between two sites.	158
8.5. An artificial layout of forty test sites, arranged in four groups of ten.	163
8.6. The results of principal component analysis of a comparison of the MDN model on the artificial four groups of 10 sites.	165
8.7. Hierarchical clustering of a comparison of the MDN model on the artificial four groups of 10 sites.	166
8.8. Scatter plot of the first two components of the vertex vectors for seven models using test data sites.	168
8.9. The similarity between seven edge models producing the same average weighted distance of $\Delta=57$ km.	169
9.1. Example of a population organized in a random network.	178
9.2. Spreading probability and final cascade size in the Watts model with $N=10000$ and $\phi_i=\phi=0.16$.	180
9.3. Spreading probabilities and corresponding average cascade size of the high-status and low-status subpopulations, when the innovation is introduced by a high-status and by a low-status individual.	183
9.4. Fractions of the high-status and low-status subpopulations that have adopted the innovation at different time steps.	184
9.5. Relationship between the parameter J and the network correlation.	187
9.6. Spreading probabilities and corresponding average cascade size of the high-status and low-status subpopulations, when the innovation is introduced by a high-status individual.	188
9.7. Spreading probabilities and corresponding average cascade size of the high-status and low-status subpopulations, when the innovation is introduced by a low-status individual.	189

9.8. Fractions of the high-status and low-status subpopulations and corresponding fractions of whole population which have adopted the innovation at different time steps.	190
9.9. Spreading probabilities for the high-status and low-status subpopulations for varying the measure of network correlation.	190

List of Tables

4.1.	Standard deviations of correlations between centrality scores in full network and sub-samples of various sampling fractions.	77
5.1.	Main characteristics of the six networks at a glance.	91
5.2.	Actors ranked according to degree score.	91
5.3.	The name of the respective network, the total number of actors, the number of influential actors, and the counts and percentages of the respective matches for each algorithm.	93
5.4.	Matches between influential actors among the top 20% highest scoring actors of the respective networks.	95
6.1.	The performance of the three algorithms that estimate transaction paths, given the set of thirteen queried paths.	114
8.1.	Physical distance measures, using quickest route or the direct distance.	152
8.2.	The ranked distances and the intervening opportunities distance.	153
8.3.	A selection of edge models which give spatial networks.	155
8.4.	The site vectors for four sites for PageRank and ‘influence betweenness’ measurements.	161
8.5.	Pearson correlation coefficients for some of the MDN networks produced by the test sites.	163
8.6.	Illustration of average weighted neighbour distance for Knossos in the four site example and using the effective distance of the quickest routes.	169

List of Contributors

Tom Brughmans is a postdoctoral researcher at the Department of Computer and Information Science, University of Konstanz. He received his PhD in Archaeology from the University of Southampton (2014) for his work entitled *Evaluating network science in archaeology. A Roman archaeology perspective*; he also gained his MSc from Southampton, in Archaeological Computing: Spatial Technologies, where he was a member of the Archaeological Computing Research Group. Tom holds an MA and BA in Archaeology from the University of Leuven. His research interests include computer applications and quantitative techniques in archaeology, digital scholarship in the humanities, landscape archaeology, and Roman archaeology. He is particularly passionate about exploring the potential of network science for the study of the past. Together with Anna Collar and Fiona Coward, Tom is a founding member of The Connected Past, a multi-disciplinary research group that aims to provide discussion platforms for the development of original and critical applications of network and complexity approaches to archaeology and history.

Fabio Caccioli is a lecturer in the Department of Computer Science at University College London. He was previously a research associate in the Centre for Risk Studies, University of Cambridge, and a postdoctoral fellow at the Santa Fe Institute (Santa Fe, USA). Fabio holds a PhD in Statistical Physics from the International School for Advanced Studies (Trieste, Italy). His research mainly focuses on systemic risk and financial stability, as well as complex networks and statistical mechanics.

Jeffery J. Clark has been a preservation archaeologist at Archaeology Southwest, a non-profit organization in Tucson, Arizona, for fifteen years. He is also an adjunct Associate Professor in the School of Anthropology, University of Arizona. Dr Clark received his PhD (1997) and MA (1990) from the University of Arizona and his BA (1983) from Cornell University. His research has focused on assessing the scale and impact of human migration, especially the Salado Phenomenon in the late pre-contact US Southwest. He has spent the past twenty-five years directing projects in southern Arizona and southwestern New Mexico as a Principal Investigator for Archaeology Southwest and a Project Director for Desert Archaeology, Inc., the latter a contract firm based in Tucson. Dr Clark has also worked extensively in Southwest Asia, including Israel, Syria, and Iraq.

Anna Collar is Assistant Professor in Classical Archaeology at Aarhus University, Denmark, working on the Sapere Aude funded project looking at the Emergence of Sacred Travel in the Graeco-Roman World. Her particular area of interest is religion in the ancient Mediterranean, and she is currently working on a book exploring the interface between Syrian migration and eastern religious innovations in the Roman world. Her first book, *Religious Networks in the Roman Empire: the Spread of New Ideas*, published by Cambridge University Press in 2013, was a finalist in the American Academy of Religion's Best New Book in the History of Religions 2014 award. Together with Fiona Coward and Tom Brughmans, Anna is a founding member of The Connected Past, a multi-disciplinary research group that aims to provide discussion platforms for the development of original and critical applications of network and complexity approaches to archaeology and history.

Fiona Coward is Lecturer in Archaeological Science at Bournemouth University, and programme co-ordinator for the BA (Hons) Archaeology and Anthropology and BSc (Hons) Anthropology degrees. Her research into the role of material culture in the evolution of the human brain and sociality has been published in a variety of journals including *Science*, the *Journal of Archaeological Science*, and the *Cambridge Archaeological Journal*, as well as in several edited volumes including *Landscapes in Mind: settlement, society and cognition in human evolution*, published by Cambridge University Press in 2014. She is a Fellow of the Royal Anthropological Institute and associate editor for archaeology for the *Journal of the Royal Anthropological Institute*, and is recorder for the Archaeology and Anthropology section of the British Science Festival. Together with Anna Collar and Tom Brughmans, Fiona is a founding member of The Connected Past, a multi-disciplinary research group that aims to provide discussion platforms for the development of original and critical applications of network and complexity approaches to archaeology and history.

Marten Düring is an historian with an interest in contemporary history and the potential of computational methods for historical research. His PhD thesis on covert support networks for persecuted Jews during the Holocaust introduces a relational perspective to the question of why people decided to help and got him interested in social network analysis. Marten has published on the applicability of network analysis methods in historical research, is co-organizer of the biennial workshop and conference series Historical Network Research and is responsible for a website with the same name. After positions as junior researcher at KWI Essen, eHumanities postdoc in Nijmegen and digital history postdoc at UNC Chapel Hill he is now working as researcher in the Digital Humanities Lab at Centre virtuel de la connaissance sur l'Europe in Luxembourg.

Tim Evans studied physics at Cambridge before completing his doctorate on thermal field theory at Imperial College London, supervised by Ray Rivers. His main interest is in the behaviour of systems of many parts, for both equilibrium and non-equilibrium situations. His early work was in the context of particle physics, cosmology, and condensed matter. However his current focus is on complex systems and in particular on the properties of complex networks. His theoretical work includes work on community detection in networks, that is, the properties of networks constrained by time and network models when space is the key constraint. He has looked at applications in social systems such as bibliometrics (citation patterns and related problems), cultural transmission, and archaeology. His archaeological collaboration with physicist Ray Rivers and archaeologist Carl Knappett started in 2002 as part of ISCOM, an EU-funded project involving social and physical scientists. He is a member of the steering committee for the Connected Past project and hosted their meeting in London in September 2014.

W. Randall ‘Randy’ Haas, Jr is a postdoctoral research associate at the Department of Anthropology, University of Maryland. He received his BA and MA from Northern Arizona University and PhD from the University of Arizona. His research interests include complex adaptive systems, social-network dynamics, hunter-gatherer anthropology, and emergent socioeconomic complexity. Regional foci include the US Southwest and the Andean highlands. His recent research used agent-based modelling, settlement-pattern analysis, and archaeological excavation in the Lake Titicaca Basin, Peru, to explore how hunter-gatherer mobility created contexts for emergent social complexity. His upcoming research with colleagues at UMD will use computational methods and complex-systems theory to model population dynamics associated with Neolithic transitions.

Anne Kandler is a lecturer in mathematics at City University London. Anne obtained her PhD in mathematics from the Technical University Chemnitz (Germany). She has held a Leverhulme Trust Early Career Fellowship at the Centre for the Evolution of Cultural Diversity at University College London and an Omidyar Fellowship at the Santa Fe Institute (USA). Her research focuses broadly on the mathematical modelling of processes of cultural evolution with a special emphasis on linking those theoretical models with empirical data describing observed episodes of cultural change. One of her major aims is to understand the theoretical limits of the statistical inference of underlying evolutionary processes from population level data such as observed temporal changes in the usage or occurrence frequencies of different cultural traits in a population.

Carl Knappett is Walter Graham/Homer Thompson Chair of Aegean Prehistory in the Department of Art at the University of Toronto. He specializes in

Aegean prehistory, pottery analysis, and archaeological theory. Recent publications include a 2011 monograph, *An Archaeology of Interaction: Network Perspectives on Material Culture and Society* and a 2013 edited volume, *Network Analysis in Archaeology: New Approaches to Regional Interaction*, both with Oxford University Press.

Barbara J. Mills is Professor in the School of Anthropology at the University of Arizona. She received her BA from the University of Pennsylvania and her MA and PhD from the University of New Mexico. She has conducted archaeological research throughout the US Southwest—especially on Ancestral Pueblo societies in the Zuni, Chaco, and Mogollon Rim areas—as well as in Guatemala, Kazakhstan, and Turkey. She conducted a long-term field project in the Silver Creek, Arizona area from 1993 to 2004 and is currently directing the collaborative Southwest Social Networks Project at the University of Arizona. Her research interests include archaeological method and theory, ceramic analysis, identity, migration, archaeologies of inequality, memory and materiality, colonialism, heritage preservation, and archaeological applications of network analysis. She is the author or editor of nine books, including the forthcoming *Oxford Handbook of the Archaeology of the Southwest*, and over sixty peer-reviewed journal articles and book chapters.

Matthew A. Peeples is Assistant Professor, School of Human Evolution and Social Change, Arizona State University, and preservation archaeologist at Archaeology Southwest, a non-profit research and preservation organization in Tucson, Arizona. He received his BA (2003) in anthropology from the University of Texas at Austin and MA (2006) and PhD (2011) degrees at Arizona State University. He has conducted archaeological research in many parts of the US Southwest including the Zuni/Cibola and Mimbres regions. His research interests are largely comparative and regional in scope, focusing on issues of social identity, social movements, and periods of rapid social change. He relies on a broad array of methods and theoretical frameworks drawing on archaeology and many other fields within the comparative social sciences.

John M. Roberts, Jr is Professor of Sociology at the University of Wisconsin-Milwaukee. His interests include social networks, quantitative methods, mathematical sociology, and organizations. He has collaborated with other participants in the Southwest Social Networks Project on a chapter in the edited volume *Network Analysis in Archaeology: New Approaches to Regional Interaction*, as well as papers in the *Journal of Archaeological Science*, the *Proceedings of the National Academy of Sciences*, and *American Antiquity*.

Ray Rivers is Emeritus Professor in Theoretical Physics and a Distinguished Research Fellow at Imperial College London. After a first degree in Mathematics at Cambridge he moved to London for a PhD on particle symmetries

with Abdus Salam's group at both Imperial College and the International Centre for Theoretical Physics, Trieste. After a subsequent postdoctoral position at the University of Chicago he returned to Imperial College London where he has been ever since, with particular emphasis on understanding quantum physics in extreme environments, such as the very early universe. However, in the last decade he has developed a parallel career in modelling social interactions, particularly those of Bronze Age Mediterranean maritime networks. This interest began with his involvement in the EU-funded ISCOM network (the Information Society as a Complex System) modelling innovation, which inevitably addressed issues of material culture including the archaeological record. In this he works closely with Tim Evans (Imperial College London) and Carl Knappett (Toronto). Ray's interests are strongly epistemological, with an aversion to using mathematical modelling in too simplistic a way.

Constantinos Tsirogiannis is a researcher in computer science, specializing in the field of algorithm design. He received his bachelor's and master's degrees in Computer Science at the National University of Athens. In 2011 he received his PhD from the Technical University of Eindhoven for his research focused on designing and implementing efficient algorithms for analysing digital representations of landscapes. Since 2011, he has been a postdoctoral researcher at the Center for Massive Data Algorithmics at Aarhus University, Denmark. His research involves the design of algorithms for applications that come from several scientific fields such as computational geometry, geographic information science, bioinformatics, and landscape archaeology.

Christos Tsirogiannis is a Greek forensic archaeologist. He studied archaeology and history of art at the University of Athens, and worked for the Greek Ministry of Culture from 1994 to 2008, excavating throughout Greece and recording antiquities in private hands. He worked with the Greek police Art Squad on a daily basis (173 cases, August 2004–December 2008) and was a member of the Greek Task Force Team that repatriated looted, smuggled and stolen antiquities from, among others, the Getty Museum, Shelby White/ Leon Levy collection, and Jean-David Cahn gallery. Since 2007, he has been working to identify in museums, galleries, auction houses, and private collections antiquities which are depicted in confiscated archives, and notifying state authorities, thus allowing them to repatriate their plundered cultural heritage. He received his PhD on the international illicit antiquities network from the University of Cambridge in 2013, when he also received the annual award for Art Protection and Security by the Association for Research into Crimes against Art (ARCA). Since May 2014, he has worked as research assistant at the Scottish Centre for Crime and Justice Research (Trafficking Culture Project), University of Glasgow.

Astrid Van Oyen is a research fellow at Homerton College, University of Cambridge. Having studied archaeology at the Universities of Leuven (Belgium) and Aix-Marseille I (France), she completed her PhD research at Cambridge on ‘Rethinking terra sigillata: an archaeological application of Actor-Network Theory’. She is currently developing various projects at the intersection of material culture studies and Roman archaeology, including research on the materiality of storage practices in the Roman Empire, fieldwork questioning craft innovation in rural Roman Italy, and a book on the material culture of the Roman Empire.

Part I

Challenging Network Methods and Theories

1

Network Perspectives on the Past: Tackling the Challenges

Tom Brughmans, Anna Collar, and Fiona Coward

INTRODUCTION

As his keynote address to the 1990 Sunbelt Social Networks conference, Mark Granovetter presented a paper entitled ‘The Myth of Social Network Analysis as a Special Method in the Social Sciences’ (Granovetter 1990). In it, he described how the **popular social network theory he proposed**, ‘The Strength of Weak Ties’ (Granovetter 1973), was like a spectre that haunted his academic career: although he subsequently pursued other research interests, he found that ‘as I got more deeply into any subject, network ideas kept coming in the back door’. He concluded that **social network analysis (SNA) is not a ‘special’ method in social science**, because ‘no part of social life can be properly analysed without seeing how it is fundamentally embedded in networks of social relations’ (Granovetter 1990: 15). However, he noted that to many, SNA is an alien concept: ‘we need to remember that there are many scholars outside the house of social network analysis who think in a relational way but don’t see the **kinship with network methods and ideas**’ (Granovetter 1990: 15).

This observation echoes the current position of network studies in archaeology and history. Few would argue that relationships between social entities are not important for understanding past social processes. However, more explicit application of network theories and methods is not yet a mainstream part of our disciplines. Although it is the case that some researchers are not aware of the advantages such perspectives might offer, the current ‘niche’ status of network applications in archaeological and historical research relates to a more general misperception: that network concepts and methodologies per se are simply not appropriate for use in research in these disciplines. This volume aims to address both issues: the contributions in this volume demonstrate both the enormous potential of network methodologies, and also—and perhaps more importantly—acknowledge and address a range of

perceived problems and reservations relating to the application of network perspectives to the study of the past, thereby encouraging and enabling their wider use in archaeology and history.

The full diversity of network perspectives has only been introduced in our disciplines relatively recently. As a result, we are still in the long-term process of evaluating which theories and methods are available, but also, more practically, the ‘fit’ between particular network perspectives and particular research questions. Many recent publications using a network perspective have showcased a range of different techniques and case studies which, while inspiring the development of the field, have almost all concluded with a note of caution: that the application of formal network perspectives to the study of the past presents a range of significant challenges which continue to hinder the broader adoption and acceptance of these methods. This volume aims to build positively on this growing maturity, to grapple with specific issues and explore potential solutions.

Our starting point is summed up by the following three observations: (1) Relational concepts are commonly used in the study of the past. However, more formal network perspectives are not necessarily the best approach in every research context. (2) Network perspectives are distinct from other research perspectives and have the potential to make a unique contribution to archaeology and history. (3) In order for this potential to be achieved a number of challenges need to be addressed.

In this opening discussion we address these observations in turn, setting the stage for the in-depth examinations of some of these challenges provided by the individual contributions that follow.

NETWORKS ARE EVERYWHERE AND NOWHERE

The use of relational concepts and relational thinking in the study of the past is not uncommon. Archaeologists and historians alike frequently discuss past networks, the interactions of entities, connectivity and the role entities fulfil as parts of a network. However, when a ‘network perspective’ is ill-defined, we might be tempted to see networks everywhere in our study of the past, without understanding the advantages this offers over other approaches. The following three excerpts illustrate the point:

In the process of diffusion and creation the isle of Crete played a foremost rôle [*sic!*]. Its geographical position enabled the Cretans to take advantage of advances made in the South and East without becoming dependent either on Egypt or on Sumer. At the same time the limited resources of their homeland obliged islanders to turn to maritime trade and thereby to diffuse their civilization along the coasts of the Mediterranean and the Black Sea. (Childe 1925: 24)

The practice of fortification can give a more reliable indication of inter-village or inter-regional hostilities. Nea Nikomedeia may conceivably have had some defensive works in early neolithic times (Rodden 1965, 84), while during the middle neolithic period two sites in Thessaly, Sesklo (Tsountas 1908, 75) and Magula Hadzimissiotiki, had enclosing walls possibly defensive in purpose, although this is not certain. Servia in Macedonia at about this time was surrounded by a ditch, possibly defensive in purpose (Heurtley 1939, 45). The outstanding example of a site apparently fortified during the neolithic period, although in the late phase, is afforded by Dhimini (dig. 18.13, 1). While less impressive in appearance on the ground than it appears in plan, the site has a central area enclosed by several lines of concentric circuit walls of stone. It is clearly the precursor of such early bronze age citadels as Troy. The concentric arrangement emphasising the significance of the central buildings contrasts markedly with the altogether democratic arrangement of the earlier neolithic villages. (Renfrew 1972: 392).

A model of the interior of a house from Ovcarovo provides an exact representation of the interior of the houses excavated at that site (Todorova 1978). I will discuss the nature of these models later in this chapter. For the moment, it is of interest to note that the 'oven' in the house model has a sloping gabled roof and openings in the front and side walls. Indeed this 'oven' looks like a 'house' model. 'The shape of the ovens is so much like that of the houses, that during excavations models of ovens have sometimes been mistaken for models of dwellings' (*ibid.*, 52). Certainly many house models do have sloping roofs, but it is unclear whether they should be reinterpreted as ovens, or whether the ovens mimic houses, or whether the Ovcarovo house model oven should be reinterpreted as a 'shrine'. Whatever the answer to this problem, it seems likely that the centrality of the oven within the house was sometimes reaffirmed by drawing symbolic parallels between houses and ovens. (Hodder 1990: 57–59)

These excerpts were drawn from sections of the books where the archaeological record is described and interpreted, rather than sections setting out the theoretical or methodological frameworks employed. Indeed, almost any other excerpt from these sections that includes some form of interpretative statement (or from books written by other archaeologists for that matter) would have served our current purpose, since the point we wish to make here is very simple: archaeologists (like historians) are interested in past phenomena that can often be abstracted using relational concepts, and they make assumptions about what entities and relationships mean, what kind of behaviour they allowed for and the implications for understanding the past phenomena they are interested in.

Childe aims to understand the role played by the island of Crete and its inhabitants in the broader Mediterranean/Black Sea region and assumes that the geographical structuring of entities allowed for and determined the spread of innovations, the creation of power dependencies, and the flow of material resources. Renfrew is interested in inter-village and inter-regional hostilities

and how this is expressed in architecture and settlement planning. He assumes that changes in the nature of hostilities gave rise to fortifications, which in turn changed the nature of those hostilities themselves. Hodder is interested in the role of house and oven models, and the symbolic parallels between them, and assumes that conceptual analogy is expressed through morphological similarities, allowing him to make interpretations regarding the centrality of the oven within the house.

All three examples explore complex past phenomena defined by multiple interacting entities (e.g. the geographical arrangements, trade routes, and cultural diffusions between Crete, Egypt, and Sumer; the spatial arrangements of regions, villages, and individual buildings; and the symbolic parallels and analogies between houses, ovens, and people) through a more abstract framework of concepts. Entities may or may not be conceptualized as clearly bounded and separable physical things, but they are regarded as analytically distinct for the purposes of the specific arguments being forwarded in each case. Network concepts can be intuitively appealing and particularly suitable for describing many of the past phenomena we aim to understand. However, to conflate 'relational' studies like this with 'network perspectives' in the more formal sense risks devaluing *both* approaches: if almost every past phenomenon can be described using loosely defined 'network' concepts we end up seeing networks everywhere, but they will rarely lead to critical insights since it is not clear what advantages they offer over other approaches. The more formal network perspectives explored in this volume come with additional criteria in terms of data requirements and/or method that make them more suitable for addressing certain research questions over others. These criteria need to be understood if formal network perspectives are to fulfil their potential of contributing new perspectives on the past.

THE POTENTIAL OF NETWORK PERSPECTIVES

To argue that network perspectives have the potential to make a unique contribution to archaeology and history implies that we believe network perspectives are a different and discrete category of research approaches. However, network perspectives are by no means a unified body of theories and methods. Network perspectives range from the highly quantitative to the highly qualitative; from those appropriate to the local social scale to those functioning best on the macro geographical or temporal scale; from scientific to philosophical; and from applications in contemporary groups to those focusing on past behaviour; and indeed every conceivable combination of the above. Each is perfectly valid, and the multi-vocality achieved by comparing results derived from different approaches has already proven extremely

fruitful for the study of the past. Here we simply wish to make two points about this diversity: first, this multiplicity of perspectives runs the risk of fostering the false impression among many non-specialists that network theories are somehow fundamentally different from network methods, and hence that network ‘thinking’ and network ‘doing’ can be easily separated. One result of this way of thinking is that specific network concepts, theories, or methods have often been applied in isolation and with little discussion of the implications of doing so. Our second point is that different network perspectives should be seen as different tools, each of which functions according to different rules. If these exciting new approaches are ever to achieve their full potential for archaeology and history, we must critically explore each one individually (for a discussion of many network theories and methods and their application to the study of the past, see Knappett 2011).

That said, and despite the diversity of network perspectives in use, they nevertheless share some common features which set them apart from other approaches and that are the key to their true potential for the study of the past. First, taking a network perspective means that the individual entities of research interest, such as technological innovations, objects, individual humans or communities, archaeological sites, and islands, are never studied in isolation. Instead, it is assumed that these entities are engaged in relationships that are fundamental to understanding their behaviour in the past. The physical size or materiality of the entities under study is largely irrelevant: almost anything can be usefully considered a node depending on the research question, potentially allowing network perspectives to bridge different spatial, social, and conceptual scales of analysis (Knappett 2011). Second, the relationships between entities can be equally diverse: a recorded action of information transmission; spatial proximity; a physical connection such as a road; friendship; political alliance; membership of an institution; presence of similar structures on different sites; the morphological similarity of objects. In this way, network perspectives allow for the incorporation of multiple entities and relationships within a single research framework. Network perspectives aim both to trace patterns of relationships between entities and to explore the implications of those relationships.

In order to do this, network perspectives require both entities and the ways in which they relate to be readily definable (even if only hypothetically or temporarily). Simplifying assumptions like this are at the heart of every theory and method applied to complex real-world phenomena, and this assumption does not mean that concepts with fluid boundaries cannot exist, that concepts cannot have multiple meanings, or that concepts cannot change their meaning, performance, or nature through time.

Any network analysis of archaeological data always implies reference to a theoretical network perspective and a level of simplification of the past phenomena studied. As we explore comprehensively in the introduction to a




Figure 1.1. Top: an abstract representation of a network model (adapted from Brandes et al. 2013, Fig. 1). Every network perspective for the study of the past includes these elements and processes. Bottom: an example of how a network model is used to explore a particular phenomenon.

recent special issue of the *Journal of Archaeological Method and Theory* (Collar et al. 2015), the key stages in the process of using a network perspective to analyse our past phenomena need to be clearly defined: that is, the processes of *abstraction* in terms of network concepts, and *representation* as network data (Brandes et al. 2013). This process is illustrated by Figure 1.1.

Archaeologists and historians aim to understand past phenomena, whether they are past networks of some sort that are hypothesized to have existed (e.g. a road network) or aspects of human behaviour that translate less straightforwardly into network concepts (e.g. trade). In Figure 1.1, the past phenomenon we are interested in is clearly separated from the perspective we use to understand it. This highlights an epistemological issue every research perspective struggles with and that is particularly critical for network perspectives as it uses intuitively appealing concepts for the study of the past: the past networks we are interested in should not be confused with the network perspective, concepts, and data we use to understand them (Isaksen 2013; Knox et al. 2006; Riles 2001). Instead of drawing a one-to-one relationship between the past phenomena and network data representations, the network perspective requires us to go through a process of abstraction in terms of network concepts and a process of representation as network data. But how exactly do we represent entities and relationships as network data, and what makes network data different from other data types?

Figure 1.2 illustrates some of these principles from a sociological perspective, but the principles can equally be applied to non-human subjects—sites, artefacts, assemblages, etc. It shows the abstraction of individuals as ‘nodes’, with a select number of attributes that the researcher has deemed relevant (in this case, gender, name, profession, income etc.: see table in Fig. 1.2). Even when we consider these individuals on their own, it is already clear that some attributes depend on others. For example, Mary’s salary is dependent on her job as an archaeologist (and potentially also on her gender). Different types of




Figure 1.2. Different data types representing individuals and social relationships: (a) individuals in isolation, characterized by their attributes which can be listed as a data table; (b) dyadic data, pairs of individuals are characterized by attributes that can only be understood with reference to both individuals (e.g. a romantic relationship); (c) network data, representing friendship relationships in this case. When one formulates the dependence assumption that a pair of individuals who have a common friend might become friends in the future, then (d) can be seen as a future development of (c); (e) network data, showing Mary, John, and James in their wider social context, comprised of relationships governed by the same dependency assumption as in (c-d).

dependencies exist when representing the relationships between individuals as edges in dyadic data (Fig. 1.2b) and network data (Fig. 1.2c). Dyadic data represents two nodes with a connection between them, and assumes that some *pairs* of nodes, rather than nodes in isolation, may also possess attributes. Being a romantically attached couple, for example (Fig. 1.2b), is an attribute that is not reducible to the attributes of either individual alone. In network data—where there are more than two nodes involved (Fig. 1.2c, here we assume all ties are friendship relationships)—we make the additional dependence assumption that one relationship will affect the existence of others. For example, one might assume that two unconnected individuals

(e.g. Mary and James) who have a friend in common (John) are more likely to become friends themselves at some point in the future (as in Fig. 1.2d). Furthermore, these three individuals are part of a wider community whose relationships might also be governed by the same dependence assumption (Fig. 1.2e).

These assumptions of dependence inherent in network data are thus fundamental for understanding the distinctive nature and potential of network perspectives. Network data requires us *to explicitly express our assumptions* about connections between nodes: what these relationships *mean*, how they affect other relationships, and what kinds of *behaviours* and *opportunities* they allow for. Network perspectives, therefore, can never take place in an ‘interpretative vacuum’: all network techniques involve the formulation of theoretical dependence assumptions.

CHALLENGES

In recent years, a growing number of researchers have begun to explore the possibilities that network perspectives offer for the study of the past. Increasing familiarity and confidence with techniques has naturally resulted in a growing appreciation of the long list of challenges faced by our disciplines in applying network perspectives (see Brughmans 2013, 2014; Isaksen 2013; Knappett 2011, 2013; Lemercier 2012; Wetherell 1998).

Some of the major challenges mentioned in previous publications and drawn from our own experiences include (loosely termed) issues of method, of data, of space and time, and of interpretation.

Method Challenges

A primary stumbling block is simply that most archaeologists and historians lack the knowledge or confidence to *critique or tailor the mathematics* behind network techniques for specifically archaeological and historical purposes (Isaksen 2013; Wetherell 1998). This can lead to the *mechanized application of network methods*, techniques and more general concepts ‘borrowed’ from other disciplines, without fully considering their appropriateness or without attempting to adapt them to the specific contexts under investigation. The theoretical underpinnings and the requirements of formal applications need to be understood in order to gauge how useful or suitable the method is in a given context (see in this volume Düring 2016; Evans 2016; Peeples et al. 2016; Van Oyen 2016). The related problem of *applying formal network perspectives to inappropriate data*, meanwhile, has the effect of devaluing the underlying

concepts and reducing their efficacy. These problems can also make the visualization stage of analysis problematic: *poor understanding of principles of network visualization* techniques result in mechanized application of layout algorithms focusing on a limited set of the network's structural features, impeding visual interpretation (particularly by non-specialists) and often reducing dynamic networks to static 'spaghetti monster' diagrams.

Data Challenges

One of the data issues arises from the fact that archaeologists and historians have borrowed many network analysis techniques from sociology. Whereas SNA focuses on *person–person* relationships, by contrast, archaeologists and historians have additionally focused on other kinds of data, for example, material or textual. An inherent *bias towards person–person relationships* risks neglecting potentially important *object–person* and *object–object* relationships (see in this volume Knappett 2016; Peeples et al. 2016; Van Oyen 2016). However, *too rigid an a priori definition of categories* of entities, attributes, hierarchical relationships, etc. is also problematic, as many are not clearly bounded, can have multiple meanings, change over time, or are appropriate only at particular scales of analysis (see in this volume Peeples et al. 2016; Van Oyen 2016). Finally, archaeologists and historians face significant problems relating to the *fragmentary nature of their datasets and the sampling biases involved* in their creation and may lack the understanding of how to formally recognize and compensate for this in network terms (see in this volume Düring 2016; Peeples et al. 2016; Tsirogiannis and Tsirogiannis 2016).

Spatio-chronological Challenges

Alongside their fragmentary nature, archaeological and historical datasets often have *poor chronological control*, which can make it difficult to explore the order in which nodes and ties emerged, and therefore the direct observation or study of processes of network change (see in this volume Peeples et al. 2016). Connected to this is the challenge posed by *issues relating to the geographical and temporal scale of analysis*. The poor resolution of much archaeological data often leads to the *overemphasis of macro-scale temporal periods and geographical regions*, whereas historical data tends to prioritize the *interpersonal scale*: valuable insights may be gained from pushing our data to work at different scales. Specifically in terms of the geographical aspects of scale, a further challenge is posed by the fact that until recently, spatial network techniques have been underdeveloped in SNA and physics,

whereas they are often significant elements of archaeological and historical research questions. This has led to limited availability of ‘off-the-shelf’ tools for exploring the differences between geographical and topological space.

Challenges in Interpretation

Finally, interpretation of the results of network analyses of archaeological and historical data is by no means unproblematic. There is a danger that the network data and concepts one uses to study past phenomena are equated with the past phenomena we are trying to understand (Isaksen 2013; Knox et al. 2006; Riles 2001). *Equifinality* also poses a potential challenge, in that many different network models can be created, each of which represent different hypothetical processes, and each of which could potentially give rise to the same observed network. Which model is best? And why? (Graham and Weingart 2015; and see in this volume Evans 2016; Kandler and Caccioli 2016; Rivers 2016).

Carl Knappett explores and critiques some of these challenges in depth in the second chapter of this volume. For Knappett, interdisciplinary dialogue and collaboration can play a crucial role in meeting these challenges—an essential step if we are to avoid the use of network perspectives becoming another ‘niche pursuit’ in the study of the past. He argues that archaeologists are particularly well-placed to bridge the apparent gaps between the more methodological and more philosophical strands of network perspectives. We further suggest that dialogue between those working on archaeological and historical datasets is particularly desirable, since both are often confronted with similar challenges and are better armed to deal with them in collaboration. Indeed, this is what the contributors to this volume have set out to demonstrate.

CONFRONTING CHALLENGES: THE STRUCTURE OF THIS BOOK

Clearly one volume cannot possibly critically address all the issues highlighted above. However, each of the contributions to this volume builds on the growing experience and confidence in applying formal network perspectives to the study of the past by mapping out and engaging in depth with one or more of these issues. Crucially, the authors of each chapter do not merely delineate the challenges, but also make constructive suggestions about how they may be overcome, while at the same time avoiding a didactic

one-size-fits-all ‘how-to’ method that would be inappropriate given the diversity of archaeological and historical datasets. While some contributions have illustrated their work with an archaeological or historical case study, these are not the central focus of this volume (for critical and innovative archaeological case studies, see papers in Knappett (2013) and *The Connected Past* special issue of the *Journal of Archaeological Method and Theory* (Collar et al. 2015)), but each of the contributions highlights issues and describes methods and techniques appropriate for all scholars of the past (and potentially for network scientists more generally).

The chapters of this book are grouped together in three parts. The first establishes the theoretical and methodological framework, while the second and third parts deal with the challenges facing researchers (re)constructing and analysing networks from historical or archaeological data and those relating to modelling dynamic network processes or evolution.

In Part 1, Knappett’s critical discussion of the many challenges from an archaeological perspective focuses on the importance of clearly defined analytical categories when using network perspectives. As he points out, this may not always be straightforward in every research context, especially where the aim is to understand the *changing* meaning attributed to entities (e.g. objects, people, ideas), when the insistence on specifying analytical categories becomes even more problematic. The following paper by Van Oyen focuses on exactly this challenge by comparing network analytical approaches with actor-network theory (ANT). Van Oyen argues that ANT ‘renders visible the heterogeneous “work-nets” needed to support and stabilize these networks and the entities they connect’, thus drawing a distinction between network analysis, which considers established categories and flows as a starting point, and ANT, which offers a ‘dynamic approach in which stability and uniformity of categories and flows is the outcome of contingent processes of maintenance work’. Explicitly differentiating between these two important but very different network concepts provides a basis for other scholars to evaluate both, with regards to their appropriateness for the distinct demands of different research contexts. Van Oyen illustrates these differences through a case study on the production of a type of Roman pottery traditionally referred to as *terra sigillata* on a single site in Central France. This case study also offers a much-needed example of how ANT can be usefully and critically adopted as a research perspective for a particular archaeological research question. Van Oyen concludes with an important cautionary statement, that the limitations of formal network analysis should be clearly understood and factored in to decisions about which methodological tools to apply in specific contexts.

The next three chapters, forming Part 2, explore issues relating to the formal analysis of networks generated from archaeological or historical data, and discuss the implications of many of the challenges highlighted in Part 1 for formal network analysis.

Peeples and colleagues identify four key features of archaeological data with important implications for the application of network methods—its material nature, variable temporal accuracy, the problem of defining network boundaries, and the impact of incomplete data. They structure their discussion via a case study on the archaeology of prehispanic settlements in the US Southwest. The authors demonstrate that these issues are inherent in much archaeological and historical data, but also that they are not insurmountable and that ‘careful consideration of research design, the methods to be used, and the assumptions involved in the creation of network datasets’ can minimize their negative impact. In particular, Peeples and colleagues emphasize the importance of clearly defining relationships between nodes: like Van Oyen, the authors stress the need to evaluate explicitly what can and what cannot be understood using different categories of relationship. They also highlight the need for careful consideration of the many assumptions and the choice of network analytical tools made during analysis in the light of the specific dataset being considered, and the impact of these decisions on the robustness of results (see also Peeples and Roberts 2013). The authors suggest hypothesis testing as a way of addressing uncertainties, and conclude with the positive observation that it is exactly those aspects of our data that historians and archaeologists usually find most challenging that hold the most promise for developing unique techniques and perspectives that will make a significant contribution to network science more broadly.

A large number of the data-related challenges introduced by Peeples and colleagues are discussed in more detail in the two chapters that follow it. Marten Düring explores one type of network measure in particular—centrality—through a historical case study focusing on support networks for persecuted Jews during the Second World War. Centrality measures are arguably the most common formal network techniques used in archaeology and history, often used to explore concepts such as power and influence. However, neither their use nor the interpretation of results are unproblematic, since these measures were often developed in other disciplines for addressing different types of research problems. Moreover, the wide range of centrality measures, while related, are all quite distinct mathematical measures, each of which therefore reveals very different aspects of a network, or of a node’s position within a network. Düring argues that in order for centrality measures to be critically and usefully applied to the study of the past, the results of individual measures and their robustness must be evaluated by applying them to the same empirical case studies and comparing results. In-depth historical analysis based on a close scrutiny of the available sources allows Düring to identify particularly influential actors in the support networks he studied, which could then be compared with those identified using a variety of different centrality measures. Centrality measures failed to identify 20–30 per cent of influential actors in this study, and so Düring concludes that although

centrality measures remain valuable in many circumstances, their potential use in studying the past requires further thought. We cannot automatically assume that measures useful to SNA will be appropriate for other kinds of datasets and research questions, or that they can be interpreted in the same way. Historians and archaeologists alike need to let their research questions and empirical datasets guide the use of such measures, and especially the interpretation of their results.

The next chapter, by Tsirogiannis and Tsirogiannis, focuses on a perennial problem in archaeological and historical research: the challenge posed by missing data. The authors tackle the issue head-on by presenting an approach that may potentially allow the identification of missing links. Their networks of illicit antiquities dealers are networks in which individuals and institutions must build the relationships of trust fundamental to these illegal transactions (similar to the support networks examined by Düring). By their very nature, such networks are very difficult to trace and will hardly ever be complete, presenting significant problems for quantitative techniques. However, Tsirogiannis and Tsirogiannis note that our best information on the complete network and the processes driving its evolution is still the observed part of the network itself: they use this information to develop five algorithms for computing transaction paths between pairs of network nodes and for indicating missing links. While these algorithms may not ‘reconstruct’ missing data to complete our networks, they do allow us to better understand our datasets and their limitations, and inform future statistical treatments of these datasets.

The final part of the volume includes three chapters dealing with the challenges facing researchers attempting to model the past dynamic processes involved in the emergence and evolution of networks.

Rivers’ chapter serves as a critical introduction to and reflection on the network modelling process for the study of the past and poses the question of whether we (intentionally or not) construct models which provide us with the results we want. In the natural sciences, models are often structured by physical laws; however, the past phenomena historians and archaeologists are traditionally interested in are not easily described by ‘laws’. Instead, the concept of individual agency is often invoked, providing considerable latitude in deciding the parameters of computational models designed for the study of the past. Rivers asks how free archaeologists and historians should be to model whatever we want, and whether simple models with few variables are often the better option. Since a variety of different models reflecting a scholar’s theoretical perspective or fitting the dataset can be constructed to explore individual agency, he asks how we can know which model is ‘right’, and how we go about comparing and evaluating competing models such as the most ‘likely’ network, or the most ‘beneficial’? In tackling these questions, Rivers focuses on exchange networks and explains some of the problems faced and decisions made in his previous work modelling maritime interaction in the

Middle Bronze Age Aegean (see Knappett et al. 2008, 2011). He concludes that in fact archaeologists and historians are rarely completely free to create whatever model we want, as the historical and archaeological research contexts as well as the available data will always provide some parameters to structure modelling efforts.

The final two chapters present more specific modelling efforts. Evans' discussion of different spatial network models follows neatly on from the paper by Rivers by putting the decision-making process suggested by Rivers into practice. Evans notes that archaeological work is often site-centric, and that the inherent nature of sites might explain why so many existing archaeological network studies are concerned with approaches to spatial clustering (e.g. Broodbank 2000; Knappett et al. 2008; Terrell 1977). While a wide range of spatial network models can be used for this purpose, Evans notes that each generates rather different kinds of results, making them difficult to interpret. Evans confronts this challenge by defining families of similar models that can be expected to produce similar results, and although he concludes that selection of the most suitable spatial network models must depend on specific research context, definition of some general 'types' of model and output allows for easier comparison and evaluation of results, and also provides a welcome introduction for archaeologists and historians interested in applying spatial network models. Indeed, Evans provides a user-friendly 'walk-through' to help scholars do this for their own research contexts.

The final chapter addresses a theme common to both history and archaeology: the spread of innovations. Archaeologists often deal with temporally aggregated data, and Kandler and Caccioli confront the challenges this presents by demonstrating how computational modelling allows the exploration of the decision-making processes of simulated individuals which may have given rise to the observed patterning. The authors present a model in which individuals possess different social statuses and are connected through social interactions that allow for the flow of information. This framework allows them to test a number of hypotheses about cultural transmission, such as the tendency for individuals of equal social status to share innovations (homophily) or the impact of the introduction of an innovation by an individual of high status. The aim of such efforts is not to prove whether a certain hypothesis is right, but rather to falsify hypotheses that are less likely to have produced the observed patterning. This chapter illustrates how computational network modelling can be a useful exercise for archaeologists and historians for thinking through their assumptions and hypotheses about a particular past phenomenon. Kandler and Caccioli's example also demonstrates that this can be done, even in the absence of archaeological and historical data, to explore the implications of specific assumptions and hypotheses and to specify the kinds of patterns one would expect to find in data that may come to light in the future.

Each of the contributions to this volume thus builds on previous work in the expanding ‘niche’ of network studies on the past to address some of the challenges early studies have encountered. Empirical examinations of some of the most common methodologies and techniques used in network analysis, such as centrality measures, establish beyond all doubt that researchers in these disciplines must become better versed in the strengths and limitations of such methods when applying them to our specific research contexts. This volume aims to arm archaeologists and historians with some of the methodological and conceptual tools they need to compare and evaluate the strengths and limitations of different approaches in specific research contexts and for particular datasets. These contributions also highlight potential and eminently practical ways forward—around, over, or straight through the problems we face when applying these methods and concepts to the study of the past. As we shake off our epistemological naivety and the field matures and becomes more sophisticated, it becomes apparent that, as well as advancing archaeological and historical research, such developments will also potentially be of huge benefit to other disciplines.

ACKNOWLEDGEMENTS

The editors would like to thank all contributors to this volume; Oxford University Press, and in particular Charlotte Loveridge and Hilary O’Shea for editorial expertise and support throughout the entire publication process; all presenters and delegates at The Connected Past conference held on 24–25 March 2012 in Southampton; Southampton’s Department of Archaeology for hosting this event, and generous funding from the following institutions that made both the conference and this publication possible: Archaeopress; the Classical Association; Oxford University Press; Southampton University Strategic Research Group Complexity in Real-World Contexts; Southampton Archaeological Computing Research Group; Southampton Web Science Doctoral Training Centre; Southampton Faculty of Humanities; and Southampton Digital Humanities.

REFERENCES

- Brandes, U., Robins, G., McCranie, A., and Wasserman, S. (2013). ‘What is Network Science?’ *Network Science*, 1:1–15.
- Broodbank, C. (2000). *An Island Archaeology of the Early Cyclades*. Cambridge: Cambridge University Press.
- Brughmans, T. (2013). ‘Thinking Through Networks: A Review of Formal Network Methods in Archaeology’, *Journal of Archaeological Method and Theory*, 20:623–62.

- Brughmans, T. (2014). 'The Roots and Shoots of Archaeological Network Analysis: A Citation Analysis and Review of the Archaeological Use of Formal Network Methods', *Archaeological Review from Cambridge*, 29:18–41.
- Childe, G. (1925). *The Dawn of European Civilization* (1st edition). London: Kegan Paul.
- Collar, A., Coward, F., Brughmans, T., and Mills, B. (2015). 'Networks in Archaeology: Phenomena, Abstraction, Representation', *Journal of Archaeological Method and Theory*, 22:1–32.
- Düring, M. (2016). 'How Reliable are Centrality Measures for Data Collected from Fragmentary and Heterogeneous Historical Sources? A Case Study', in T. Brughmans, A. Collar, and F. Coward (eds), *The Connected Past: Challenges to Network Studies in Archaeology and History*. Oxford: Oxford University Press, 85–101.
- Evans, T. (2016). 'Which Network Model Should I Use? Towards a Quantitative Comparison of Spatial Network Models in Archaeology', in T. Brughmans, A. Collar, and F. Coward (eds), *The Connected Past: Challenges to Network Studies in Archaeology and History*. Oxford: Oxford University Press, 149–73.
- Graham, S., and Weingart, S. (2015). 'The Equifinality of Archaeological Networks: An Agent-Based Exploratory Lab Approach', *Journal of Archaeological Method and Theory*, 22:248–74.
- Granovetter, M. (1973). 'The strength of weak ties', *American Journal of Sociology*, 78(6):1360–80.
- Granovetter, M. (1990). 'The Myth of Social Network Analysis as a Special Method in the Social Sciences', *Connections*, 13–16.
- Hodder, I. (1990). *The Domestication of Europe: Structure and Contingency in Neolithic Societies*. Oxford: Blackwell.
- Isaksen, L. (2013). ‘“O What A Tangled Web We Weave”—Towards a Practice That Does Not Deceive’, in C. Knappett (ed) *Network Analysis in Archaeology. New Approaches to Regional Interaction*. Oxford: Oxford University Press, 43–70.
- Kandler, A., and Caccioli, F. (2016). 'Networks, Homophily, and the Spread of Innovations', in T. Brughmans, A. Collar, and F. Coward (eds), *The Connected Past: Challenges to Network Studies in Archaeology and History*. Oxford: Oxford University Press, 175–97.
- Knappett, C. (2011). *An Archaeology of Interaction. Network Perspectives on Material Culture and Society*. Oxford and New York: Oxford University Press.
- Knappett, C. (ed.) (2013). *Network Analysis in Archaeology. New Approaches to Regional Interaction*. Oxford: Oxford University Press.
- Knappett, C. (2016). 'Networks in Archaeology: Between Scientific Method and Humanistic Metaphor', in T. Brughmans, A. Collar, and F. Coward (eds), *The Connected Past: Challenges to Network Studies in Archaeology and History*. Oxford: Oxford University Press, 21–33.
- Knappett, C., Evans, T., and Rivers, R. (2008). 'Modelling maritime interaction in the Aegean Bronze Age', *Antiquity*, 82:1009–24.
- Knappett, C., Evans, T., and Rivers, R. (2011). 'The Theran eruption and Minoan Palatial Collapse: New Interpretations Gained from Modelling the Maritime Network', *Antiquity*, 85:1008–23.

- Knox, H., Savage, M., and Harvey, P. (2006). 'Social Networks and the Study of Relations: Networks as Method, Metaphor and Form', *Economy and Society*, 35:113–40.
- Lemercier, C. (2012). 'Formal Network Methods in History: Why and How?', in G. Fertig (ed) *Österreichische Zeitschrift für Geschichtswissenschaften*. Turnhout: Brepols.
- Peeples, M. A., Mills, B. J., Haas, W. R. J., Clark, J. J., and Roberts, J. M. J. (2016). 'Analytical Challenges for the Application of Social Network Analysis in Archaeology', in T. Brughmans, A. Collar, and F. Coward (eds), *The Connected Past: Challenges to Network Studies in Archaeology and History*. Oxford: Oxford University Press, 59–84.
- Peeples, M. A., and Roberts, J. M. (2013). 'To Binarize or not to Binarize: Relational Data and the Construction of Archaeological Networks', *Journal of Archaeological Science*, 40:3001–10.
- Renfrew, C. (1972). *The Emergence of Civilisation: The Cyclades and the Aegean in the Third Millennium BC*. London: Methuen.
- Riles, A. (2001). *The Network Inside Out*. Ann Arbor, MI: University of Michigan Press.
- Rivers, R. (2016). 'Can Archaeological Models Always Fulfil our Prejudices?', in T. Brughmans, A. Collar, and F. Coward (eds), *The Connected Past: Challenges to Network Studies in Archaeology and History*. Oxford: Oxford University Press, 123–47.
- Terrell, J. E. (1977). 'Human Biogeography in the Solomon Islands', *Fieldiana Anthropology*, 68:1–47.
- Tsirogiannis, C., and Tsirogiannis, C. (2016). 'Uncovering the Hidden Routes: Algorithms for Identifying Paths and Missing Links in Trade Networks', in T. Brughmans, A. Collar, and F. Coward (eds), *The Connected Past: Challenges to Network Studies in Archaeology and History*. Oxford: Oxford University Press, 103–20.
- Van Oyen, A. (2016). 'Networks or Work-nets? Actor-Network Theory and Multiple Social Topologies in the Production of Roman *Terra Sigillata*', in T. Brughmans, A. Collar, and F. Coward (eds), *The Connected Past: Challenges to Network Studies in Archaeology and History*. Oxford: Oxford University Press, 35–56.
- Wetherell, C. (1998). 'Historical Social Network Analysis', *International Review of Social History*, 43:125–44.

2

Networks in Archaeology: Between Scientific Method and Humanistic Metaphor

Carl Knappett

INTRODUCTION: ARCHAEOLOGY AS BRIDGE BETWEEN DISCIPLINES

Running the full gamut of scholarship from physical science to philosophy, archaeology's diversity can be a negative rather than a positive—when the same phenomena can attract such different approaches that archaeologists end up talking past one another. Take the example of archaeological landscape analysis: on the one hand, this has produced rich, expressive phenomenological studies, and on the other, detailed palaeoenvironmental reconstructions. These two perspectives are not commonly combined, although when they are archaeology emerges as a solid bridge between the humanities and the environmental sciences (van der Leeuw and Redman 2002; Smith et al. 2012). Take another example: artefact studies—on the one hand, poetic, philosophical musings on anything from fabulous artworks to mundane artefacts, and on the other, neutron activation analysis, X-ray diffraction and petrography employed in characterizing stone and ceramic technologies. There are calls (e.g. Jones 2004; Sillar and Tite 2000) to 'humanise' the science (and vice versa would also be fitting), and perhaps in artefact studies the integration of the sciences and humanities has had more success than in landscape studies. It is a difficult balancing act. But those archaeological studies that do find a way to combine both often create more convincing interpretations.

Alongside landscape and artefact studies, network analysis is a third exemplar of this tension between scientific and humanistic understandings in archaeology. On the one hand, networks can be used quite formally and quantitatively to analyse interactions in space or, indeed, cultural evolution over time (Henrich and Broesch 2011). This use of networks is quite different from a more qualitative, figurative use, as seen recently in book-length

treatments by Irad Malkin (2011) and Ian Hodder (2012). There is a danger of the gap between these different understandings of networks widening, just as the humanistic and scientific understandings of both landscapes and artefacts can sometimes seem incommensurate. I think one can see a certain reticence about being sucked into the ‘scientism’ of networks, what one might even dub ‘networkitis’, along the lines of ‘Darwinitis’ or the tendency for all manner of cultural phenomena to now find ‘explanation’ through evolutionary models, and recently the subject of a stinging critique by Raymond Tallis (2011).

However, a pattern that seems real, and very promising, is that many of the recent studies using network analysis in archaeology avoid this tension. They are both humanistic and scientific. A very good example is the work of Barbara Mills and colleagues in the Southwest Social Networks Project (Mills et al. 2013a, 2013b), and one could say the same of many of the contributions to a recent edited volume, *Network Analysis in Archaeology: New Approaches to Regional Interaction* (Knappett 2013). All use network analysis in some formal sense to address archaeological questions across a broad range of case studies, from Oceania to the Mediterranean, and from the Epipalaeolithic to the medieval period. One interesting feature that emerges in the volume is a healthy critique of network approaches—perhaps a sign of approaching maturity, after the early glow of excitement. Certainly, archaeology needs to be careful about applying techniques from other fields uncritically. Søren Sindbæk warns us that archaeologists should realize that, unlike sociologists, we lack actual networks with a full set of nodes and links to analyse. Instead, we have to guess what they are (as if in a black box) from their inputs and outputs, which is, in fact, more akin, he argues, to a process of network *synthesis* (Sindbæk 2013). Leif Isaksen observes that archaeologists need to be careful how they use network metrics and visualizations when, as is often the case, they do not have a strong background in mathematics (Isaksen 2013). Erik Gjesfjeld and Colby Phillips offer a different kind of critique, that the network structures we identify archaeologically may not tally with traditional models of exchange, such as reciprocity and redistribution, meaning that archaeologists need to generate new exchange models that place greater emphasis on network relationships (Gjesfjeld and Phillips 2013).

Such critiques are important. We need not content ourselves with being laggardly, passive consumers of innovation in other disciplines. In fact, archaeology is uniquely placed—poised between the humanities, social sciences, and the hard and soft sciences—to make a distinctive contribution to methodological and philosophical bridge-building in network thinking. One need only look at the strained relationship between social network analysis (SNA) and social physics, for example, to see the problems that sometimes stem from a lack of communication across fields (e.g. Scott 2011; Watts 2004). My argument, specifically, is that archaeologists necessarily engage with some key features of networks that arguably have not been given due attention,

either from the more humanistic or the scientific sides. These can be treated under five categories: space, heterogeneity, scale, temporality, and incompleteness. I now deal with each of these in turn.

SOCIAL AND PHYSICAL SPACE

The first of these key features concerns the disjuncture between social and physical space and, in particular, the tendency to focus on the former at the expense of the latter. One source of ideas for networks in archaeology recently has been SNA. As the name suggests, its focus is social networks, in which the nodes are commonly individuals and the links their relations, such as friendships (Breiger 2004; Scott 2011). SNA aims to reveal social structure through analysis of networks of this kind, identifying cliques, different forms of centrality, or structural holes (Burt 1992). The focus is very much on social space, or topology, rather than physical, Euclidean space. In analysing friendships in a university, for example, the shared physical backdrop of the university is implied rather than foregrounded and is rarely given an explanatory role (Bidart et al. 2011: 209), though there are important exceptions, such as the work of Barry Wellman and colleagues (Wellman 1979; Mok et al. 2007) and Carter Butts (2003). In a recent paper on networks and geography, Daraganova and colleagues state that ‘there have been relatively few attempts to build explicitly spatial models of social networks and to use these models to understand the way in which social networks are embedded geographically’ (Daraganova et al. 2012: 6).

It is also the case that another domain of network analysis that has found its way into archaeology, complexity science, has been more focused on topology, on relational rather than physical space, as in the ground-breaking work on ‘small worlds’ (Watts and Strogatz 1998). Barthelemy, in a recent review, notes that spatial aspects of networks hardly feature in recent network science textbooks (2011: 3). And yet, he stresses, space and distance are certainly relevant in many networks, as was recognized some decades ago in quantitative geography, and seen quite explicitly in the work of Haggett and Chorley (Barthelemy 2011: 3). As Barthelemy establishes in his fine review, however, network scientists are now giving geometric space the attention it merits, looking particularly at the effects of cost on edge dynamics, for instance, in transportation and communication networks (e.g. Expert et al. 2011).

This is a curious inversion of what has happened in archaeology, which has had a long interest in patterning in physical space, indeed going back at least as far as the influence of New Geography (e.g. Haggett and Chorley 1969), and continuing through processual approaches and the considerable uptake of geographic information systems (GIS) in the discipline (e.g. Conolly and

Lake 2006). This ‘physical’ emphasis has more recently been added to by more of a concern for social space, often through network analysis. Arguably, the conjunction of these two elements raises particular issues for archaeologists, in that we are obliged to reconstruct social networks from physical patterning in geometric space (see Sindbæk 2013 on network synthesis). One of our fundamental problems, then, is how to effect an articulation of physical and social space. As discussed above, this is also a growing concern in both SNA and social physics, though both risk reinventing the wheel by overlooking some of the earlier work in quantitative geography (e.g. Haggett and Chorley 1969). Although this also had faded from view somewhat with the influence of post-processual approaches in the 1980s and 1990s, such geographical approaches were always present. One example is the collaboration with archaeologists of geographer Alan Wilson, who did important work in the 1970s on modelling the spatial distribution of retail outlets. He collaborated with an ancient historian to model the distribution of Greek city-states in Attica (Rihll and Wilson 1991) and has recently returned to ancient Greece in a new collaboration on the network modelling of settlement distribution in Bronze Age Crete (Bevan and Wilson 2013).

HETEROGENEITY: PEOPLE AND THINGS

A second key feature of the networks with which archaeologists are concerned is that they are not purely social networks, but sociomaterial networks. That is to say, the networks are heterogeneous, composed of both humans and nonhumans or, less opaquely, people *and* things. This is not something one typically sees in SNA, though it is an important tenet of actor–network theory (ANT). Indeed, ANT insists on analytical symmetry between human and nonhuman actants, rather than assuming from the outset the primacy of the human. In archaeology, Ian Hodder has recently developed an approach that he calls ‘entanglement’, which is quite explicit about this heterogeneity and the different kinds of relations that exist, not only between humans and humans and between humans and things (and vice versa), but also between things themselves. He has framed this as follows, with ‘H’ standing for humans and ‘T’ for things (Hodder 2012: 88):

$$\text{Entanglement} = (\text{HT}) + (\text{TT}) + (\text{TH}) + (\text{HH})$$

Hodder does make some use of the term ‘network’, but stops short of employing formal network methods. Although in his scheme ‘HT’ and ‘TH’ relations imply a need for bipartite, or two-mode, networks, that is with two different kinds of nodes, humans and things, Hodder does also imply that there are direct thing–thing relations (‘things depend on other things’), which

might presumably be depicted as a one-mode network. However, he does then qualify this by insisting that ‘humans are involved in these chains from the start. And they remain pulled into the links and interstices . . . [humans] become embroiled in the dependences and dependencies between things’ (Hodder 2012: 59). This, then, would suggest that even thing–thing relations should be conceived of as mediated by human connections. The issue then becomes whether one gains anything conceptually by inserting human ‘nodes’ to form a bipartite network—especially in prehistoric archaeology when these nodes are almost always invisible. If we cannot insert a meaningful category of ‘human’ as a node between things, then perhaps we should draw direct links between things. Those links may represent an iconic or indexical connection, which does, of course, imply an interpretative (human) act linking two nodes. In other words, the indexical, causal connections one might draw between different kinds of pot, signifying perhaps that they all derive from the same workshop, indicates a common artisan in the interstices between them. That these kinds of human–thing dependencies could be suitably modelled in network form finds support in the idea recently expressed by Brandes and colleagues that all complex networks are structured around dependencies, with one example being preferential attachment (Brandes et al. 2013: 10).

MULTIPLE SCALES

Third, archaeologists are obliged to think across scales, though we tend not to do this very explicitly. We need to be able to zoom in to the feature, the installation, the household, and then out to the community and the region. Different approaches have been employed, and they can be broadly categorized as either top-down or bottom-up. World-systems theory is perhaps the best known of the top-down approaches, and though it has met with some success at the broader, inter-regional level, it has struggled to work its way down to the micro-scale (e.g. Sherratt and Sherratt 1991; see also Kristiansen and Larsson 2005). World-systems theory has certainly met with criticism for this, and Gil Stein (2002) has underlined the importance of developing approaches to inter-regional interaction that are more sensitive to the micro-scale (see also discussion in Knappett 2011a: 27; Harding 2013). In contrast, Clive Gamble develops a bottom-up approach, initially in connection with his work on Palaeolithic societies (Gamble 1998; 1999), but actually with broader applicability to later periods too. He begins with the ‘intimate’ networks of face-to-face interactions and scales up through ever increasing levels to ‘global’ networks. The idea of a gradual ‘release from proximity’ is significant in his argument, and he sees material culture as playing a key role in this

capacity of human communities to maintain far-flung networks of interaction (see also Gamble 2007). This provides us with a neat link to the previous discussion on heterogeneity, as people–thing interactions are key in Gamble’s scheme (see also Knappett 2011a: 30).

The multi-scale approach advocated by Gamble, based around network thinking, has much to offer, particularly for the very long-term processes with which he is principally concerned. However, archaeologists continue to find other kinds of solutions to the multi-scale problem. For example, John Robb has recently rekindled discussion of scalar issues using a classic archaeological example: the spread of farming across Europe in the Neolithic (Robb 2013). On the one hand, this process is the result of a whole series of varied local decisions and practices. On the other hand, it is a macro-scale phenomenon with a high degree of convergence continent-wide. Current approaches do a poor job of integrating these different scales: explanations play out at either the micro- or the macro-scale. Robb’s solution is to devise a meta-model that can help to bring the local and continental scales together. He proposes using notions of ‘emergent causation’ and ‘landscapes of action’ to think through the changes across these apparently divergent scales. As is apparent from the discussion above, I would argue that network thinking suits this kind of multi-scalar analysis (Gamble 1999; Knappett 2011a). It is unclear how Robb’s concepts relate to network approaches, however. It may be that Robb considers networks as passive and static; certainly, his use of the words ‘emergent’ and ‘action’ suggests a concern with dynamics. This also seems to be a perception guiding such approaches in related disciplines that are also very much concerned with the problem of multiple scales, such as sociology and human geography. Mimi Sheller, for example, writes that: ‘Rather than mathematically precise network analytical approaches to describing social worlds, I argue that a “messier” imagery of liquid social dynamics will enable a better understanding of the complexity of these mobile social interactions’ (Sheller 2004: 41).

Another critique sees networks as just one dimension of sociospatial relations, alongside territories, places, and scales, and argues that a polymorphic approach is needed to understand how they are intertwined (Jessop et al. 2008). This is a valiant attempt to make sense of the many different spatial turns, and to find a coherent and methodologically sound way forward, though in their schema it remains difficult to grasp the different roles of networks, territories, scales, and places. Perhaps more interdisciplinary dialogue between archaeology and this kind of approach to sociospatial relations across scales would be beneficial. Some efforts have been made, as seen in the engagement with the work of Edward Soja, for example, and what emerges is the importance of wrapping time into sociospatial relations (Blake 2002; Swenson 2012). Archaeologists, evidently, have a particular concern with considering space from a ‘deep history’ perspective (Shryock and Smail 2011).

TEMPORALITY

We can use this concern with time as a segue into our fourth point, the importance of long-term processes to archaeology, particularly as we have briefly discussed the spread of farming, a process that in Europe pans out over three millennia. The salience of network thinking and analysis might, however, seem rather tenuous, as evolution has been something of a sore point in much network analysis, with networks easily interpreted as atemporal and ahistoric. Butts (2009) has recently underlined the importance of considering the time scale at which the phenomenon under study plays out. He argues that while sometimes it may be appropriate to depict a static network, in other instances this can dramatically skew the results. He provides two contrasting examples. On the one hand, in order to understand rapid processes of information diffusion through a network, friendships are probably sufficiently stable to justify static representation. On the other hand, representing sexual partnerships as a static network when trying to understand the epidemiology of sexually transmitted diseases could be problematic if those partnerships alter more quickly than the disease spreads. Studying time evolution in networks does present some technical challenges, but network scientists are devoting more and more attention to techniques for representing dynamic processes (see Holme and Saramäki 2012; also Nicosia et al. 2013). One method is to take a series of freeze-frames over the lifetime of a network and to stitch them together like a flipbook—a so-called ‘multi-slice network’ (Mucha et al. 2010). Comparing a series of steady network states over time is essentially also what Emma Blake does in her archaeological network analysis of the emergence of ethnic groups in the prelude to the Roman Empire (Blake 2013). She employs the idea of ‘path dependence’ to examine the extent to which later ethnic groups in the Italian peninsula actually take shape as social networks much earlier, as far back as the Late Bronze Age.

However, these examples deal with preordained categories—whether they are individuals or sites—which network analysis then joins together to allow a better understanding of their patterns of connection. But what if it is the very emergence of a category that one wishes to understand? Astrid Van Oyen has made the important point that whereas SNA has tended to work from categories in this way, ANT has stressed the need to analyse the ‘work-nets’ that lead to the very stabilization of categories in the first place (Van Oyen 2016). Van Oyen’s example is the widespread Roman pottery ware *sigillata*. SNA would simply take *sigillata* as given, and proceed to analyse the connections between different *sigillata* workshops or consumers, to find, for example, patterns of centrality. ANT, on the other hand, would look at the various processes that serve to make *sigillata* into a stable category, and the ongoing investments (work-nets) that maintain it as such. This certainly entails a

concern for temporality, but it is rather different to that expressed in the ‘timeslice’ approach outlined above. Instead, it is about issues of the emergence of meaning and is consistent with other approaches that are gaining ground in archaeology, such as Ian Hodder’s ‘entanglement’ approach, mentioned above (2012).

Hodder does begin by talking in terms of networks, but ultimately eschews them for their inability to convey the ‘stickiness’ of human–thing dependencies (Hodder 2012: 94). ‘Dependency’—implying temporality (see also ‘path dependence’ above, Blake 2013)—is also of principal concern in Tim Ingold’s ‘ecology of materials’ (2012). For Ingold, ‘materials do not exist as static entities with diagnostic attributes . . . [materials are histories] . . . materials, thus, *carry on*, undergoing continual modulation as they do so’ (Ingold 2012: 434–5). Thus, he argues against point-to-point thinking, stressing that we must instead follow the flow, and that the craftsman must try to couple his or her own movements and gestures with those of the material. However, this is not a connection between entities, but an entangling of flows, and as such Ingold argues that the term ‘meshwork’ better captures this ‘web of life’ than does ‘network’. There is evidently a pattern here, a reluctance to embrace network thinking because of some lingering doubt about the capacity of networks to cope with temporality and change. I have argued elsewhere that it is quite possible to work between both meshworks and networks, as they offer different kinds of emphasis (Knappett 2011b). Many of the misconceptions about their perceived incompatibilities are less to do with network methods *per se*, and more with the static ways in which they have often been used. Much more work is needed to render network methods more dynamic and sensitive to time.

INCOMPLETENESS

Fifth and finally, we have incomplete data. Although usually regarded as a weakness, incompleteness provides interesting challenges for modelling networks. If we are attempting to model the dynamics of settlement interactions, how can we be sure that our results are not severely compromised by the inevitable gaps in our knowledge? In other words, we cannot be sure that all sites in a given period and region are known archaeologically, so what difference would the addition of a new site make to our results? One solution to this problem of robustness is what we might call ‘coarse-graining’ (Evans et al. 2009: 460–1). That is to say, a model is constructed in such a way that the lower-level details do not affect the macro-scale patterning. So, for example, a model might include the assumption that detailed interactions at the local

level of a particular island are less important than the wider interactions between islands. If this is the case, then one can effectively aggregate both sites and interactions at the island level, which is to adopt a ‘centre of mass’ approach (Evans et al. 2009: 461). This in turn means that the model ought to be gravitational, ‘since the attribute of gravitational energy is that it is the same, whether it is calculated from the centres of mass, or from the individual constituents of those masses’ (Evans et al. 2009: 461). Therefore a gravity model, in this sense, minimizes the effects of the partial and fragmentary nature of the archaeological record.

Another, similar solution to incompleteness has recently been proposed by Bevan and Wilson (2013). As with Evans et al. (2009), their focus is on settlement patterns in the Bronze Age Aegean, though aimed specifically at the island of Crete. They too bring up gravity models as a possible, partial solution, though phrased in terms of entropy maximization, and connected with Wilson’s earlier work on the distribution of retail outlets. However, they note that while such models help to predict the relative size of known locations, in archaeology it is rare to have a complete record of settlement locales. To actually propose the likely location of missing settlements based on existing evidence, they suggest inhomogeneous point process modelling (Bevan and Wilson 2013: 2416). Devising such a model for Bronze Age Crete, they link the inhomogeneous distribution of sites to the uneven distribution of agricultural land; by also including likely interactions along the shortest paths between sites and including a measure for the higher importance of coastal sites (due to inflow of connections from off-island), their results do show quite well the likely locations of ‘missing’ sites, such as in parts of west Crete that have not been as intensively researched as the centre and east of the island. Including such unknown sites in network models of island interactions would then seem reasonable, given the measures taken to deal with the inevitable incompleteness of the evidence.

Sindbæk (2013) takes a rather different perspective, arguing that the incompleteness of the archaeological record means that we usually cannot do genuine network analysis. That is to say, one requires all the nodes and links, as is usually the case in SNA, in order to be able to analyse them. In archaeology, we are invariably faced with ‘the task of reconstructing the broken links of a ruined network from observable distributions and patterns of association’ (Sindbæk 2013: 71). He then suggests that this makes the task one of network *synthesis* rather than analysis: it is akin to reconstructing black-box circuits, when given only inputs and outputs to work with. This being the nature of the problem, Sindbæk goes on to argue that network design optimization, which deals with exactly this kind of black-box problem, would be a more appropriate source of inspiration for archaeological networks than the commonly used SNA.

CONCLUSIONS

Each of these five issues—space, heterogeneity, scale, temporality, and incompleteness—is important in both network science and the humanities. With greater dialogue between disciplines to better grasp these issues—and particularly, I argue, a stronger input from archaeology—we can hopefully move towards a more satisfactory application of network ideas. However, this is not a totalizing call for network methods above all else. The network is, after all, a form of abstraction, of ordering, and some ancient phenomena, at whatever scale, may simply not be of a topology that fits. Still, network analysis can in itself reveal this—when, for example, an ancient technology just does not quite fit into a *chaîne opératoire* diagram. More pressing than a broader use of network techniques is, I think, fuller attention to relations and interactions. There are different paths towards this goal, though I would argue that networks are a particularly good solution: they can transform our efforts to think relationally, as ANT has shown. And we should use their full potential, which goes far beyond their use at the regional scale for understanding exchange. They can be incredibly useful across all scales, from the micro- to the macro-level. And an even more radical and challenging use of network thinking is for exploring human–artefact entanglements, something that we can do over the very long term to create deep histories. Networks have enormous potential to bring together interests from the different sides of archaeology, from the philosophical to the physical sciences. If we do not attempt to harness the full power of networks, this incredibly useful theory and method may slip away and become little more than a niche pursuit—and archaeology has enough of these already.

ACKNOWLEDGEMENTS

I am grateful to Anna Collar, Fiona Coward, and Tom Brughmans for their invitation to the conference and for their useful feedback on an earlier draft.

REFERENCES

- Barthelemy, M. (2011). ‘Spatial Networks’, *Physics Reports*, 499:1–101.
- Bevan, A., and Wilson, A. (2013). ‘Models of Settlement Hierarchy Based on Partial Evidence’, *Journal of Archaeological Science* 40:2415–27.
- Bidart, C., Degenne, A., and Grossetti, M. (2011). *La Vie en Réseau: Dynamique des Relations Sociales*. Paris: Presses Universitaires de France.
- Blake, E. (2002). ‘Spatiality Past and Present: An Interview with Edward Soja, Los Angeles, 12 April 2001’, *Journal of Social Archaeology*, 2:139–58.

- Blake, E. (2013). 'Social Networks, Path Dependence and the Rise of Ethnic Groups in Pre-Roman Italy', in C. Knappett (ed), *Network Analysis in Archaeology: New Approaches to Regional Interaction*. Oxford: Oxford University Press, 203–21.
- Brandes, U., Robins, G., McCranie, A., and Wasserman, S. (2013). 'What is Network Science', *Network Science*, 1:1–15.
- Breiger, R. L. (2004). 'The Analysis of Social Networks', in M. Hardy and A. Bryman (eds), *Handbook of Data Analysis*. London: Sage, 505–26.
- Burt, R. S. (1992). *Structural Holes: The Social Structure of Competition*. Cambridge, MA: Harvard University Press.
- Butts, C. T. (2003). 'Predictability of Large-Scale Spatially Embedded Networks', in R. Breiger, K. Carley, and P. Pattison (eds), *Dynamic Social Network Modelling and Analysis: Workshop Summary and Papers*. Washington, DC: National Academies Press, 313–23.
- Butts, C. T. (2009). 'Revisiting the Foundations of Network Analysis', *Science*, 325:414–16.
- Conolly, J., and Lake, M. (2006). *Geographical Information Systems in Archaeology*. Cambridge: Cambridge University Press.
- Daraganova, G., Pattison, P., Koskinen, K., Mitchell, B., Bill, A., Watts, M., and Baum, S. (2012). 'Networks and Geography: Modelling Community Network Structures as the Outcome of Both Spatial and Network Processes', *Social Networks*, 34:6–17.
- Evans, T., Knappett, C., and Rivers, R. (2009). 'Using Statistical Physics to Understand Relational Space: A Case Study from Mediterranean Prehistory', in D. Lane, D. Pumain, S. van der Leeuw, and G. West (eds), *Complexity Perspectives on Innovation and Social Change*. Berlin: Springer, 451–80.
- Expert, P., Evans, T. S., Blondel, V. D., and Lambiotte, R. (2011). 'Uncovering Space-Independent Communities in Spatial Networks', *Proceedings of the National Academy of Sciences*, 108(19):7663–8.
- Gamble, C. (1998). 'Palaeolithic Society and the Release from Proximity: A Network Approach to Intimate Relations', *World Archaeology*, 29:426–49.
- Gamble, C. (1999). *The Palaeolithic Societies of Europe*. Cambridge: Cambridge University Press.
- Gamble, C. (2007). *Origins and Revolutions: Human Identity in Earliest Prehistory*. Cambridge: Cambridge University Press.
- Gjesfjeld, E., and Phillips, S. C. (2013). 'Evaluating Adaptive Network Strategies with Geochemical Sourcing Data: A Case Study from the Kuril Islands', in C. Knappett (ed), *Network Analysis in Archaeology: New Approaches to Regional Interaction*. Oxford: Oxford University Press, 281–305.
- Haggett, P., and Chorley, R. J. (1969). *Network Analysis in Geography*. London: Edward Arnold.
- Harding, A. F. (2013). 'World Systems, Cores, and Peripheries in Prehistoric Europe', *European Journal of Archaeology*, 16:378–400.
- Henrich, J., and Broesch, J. (2011). 'On the Nature of Cultural Transmission Networks: Evidence from Fijian Villages for Adaptive Learning Biases', *Philosophical Transactions of the Royal Society B*, 366:1139–48.
- Hodder, I. (2012). *Entangled: An Archaeology of the Relationships between Humans and Things*. Malden, MA: Wiley-Blackwell.

- Holme, P., and Saramäki, J. (2012). ‘Temporal Networks’, *Physics Reports*, 519:97–125.
- Ingold, T. (2012). ‘Toward an Ecology of Materials’, *Annual Review of Anthropology*, 41:427–42.
- Isaksen, L. (2013). ‘“O What a Tangled Web We Weave”—Towards a Practice That Does Not Deceive’, in C. Knappett (ed), *Network Analysis in Archaeology: New Approaches to Regional Interaction*. Oxford: Oxford University Press, 43–67.
- Jessop, B., Brenner, N., and Jones, M. (2008). ‘Theorizing Sociospatial Relations’, *Environment and Planning D: Society and Space*, 26:389–401.
- Jones, A. (2004). ‘Archaeometry and Materiality: Materials-Based Analysis in Theory and Practice’, *Archaeometry*, 46:327–38.
- Knappett, C. (2011a). *An Archaeology of Interaction: Network Perspectives on Material Culture and Society*. Oxford: Oxford University Press.
- Knappett, C. (2011b). ‘Networks of Objects, Meshworks of Things’, in T. Ingold (ed), *Redrawing Anthropology: Materials, Movements, Lines*. Anthropological Studies of Creativity and Perception. Farnham: Ashgate, 45–63.
- Knappett, C. (ed) (2013). *Network Analysis in Archaeology: New Approaches to Regional Interaction*. Oxford: Oxford University Press.
- Kristiansen, K., and Larsson, T. B. (2005). *The Rise of Bronze Age Society: Travels, Transmissions and Transformations*. Cambridge: Cambridge University Press.
- Malkin, I. (2011). *A Small Greek World: Networks in the Ancient Mediterranean*. Oxford: Oxford University Press.
- Mills, B. J., Clark, J. J., Peebles, M. A., Haas, W. R. Jr, Roberts, J. M. Jr, Hill, J. B., Huntley, D. L., Borck, L., Breiger, R. L., Clauset, A., and Shackley, M. S. (2013a). ‘Transformations of Social Networks in the Late Pre-Hispanic US Southwest’, *Proceedings of the National Academy of Sciences*, 110(15):5785–90.
- Mills, B. J., Roberts, J. M. Jr, Clark, J. J., Haas, W. R. Jr, Huntley, D., Peebles, M. A., Borck, L., Ryan, S. C., Trowbridge, M., and Breiger, R. L. (2013b). ‘The Dynamics of Social Networks in the Late Prehispanic US Southwest’, in C. Knappett (ed), *Network Analysis in Archaeology: New Approaches to Regional Interaction*. Oxford: Oxford University Press, 181–202.
- Mok, D., Wellman, B., and Basu, R. (2007). ‘Did Distance Matter Before the Internet? Interpersonal Contact and Support in the 1970s’, *Social Networks*, 29:430–61.
- Mucha, P. J., Richardson, T., Macon, K., Porter, M. A., and Onnela, J.-P. (2010). ‘Community Structure in Time-Dependent, Multiscale, and Multiplex Networks’, *Science*, 328:876–8.
- Nicosia, V., Tang, J., Mascolo, C., Musolesi, M., Russo, G., and Latora, V. (2013). ‘Graph Metrics for Temporal Networks’, in P. Holme and J. Saramäki (eds), *Temporal Networks*. Berlin: Springer, 15–40.
- Rihll, T. E., and Wilson, A. G. (1991). ‘Modelling Settlement Structures in Ancient Greece: New Approaches to the Polis’, in J. Rich, and A. Wallace-Hadrill (eds), *City and Country in the Ancient World*. London: Routledge, 59–95.
- Robb, J. (2013). ‘Material Culture, Landscapes of Action and Emergent Causation: A New Model for the Origins of the European Neolithic’, *Current Anthropology*, 54:657–83.
- Scott, J. (2011). ‘Social Network Analysis: Developments, Advances and Prospects’, *Social Network Analysis and Mining*, 1:21–6.

- Sheller, M. (2004). 'Mobile Publics: Beyond the Network Perspective', *Society and Space*, 22:39–52.
- Sherratt, A. G., and Sherratt, E.S. (1991). 'From Luxuries to Commodities: The Nature of Mediterranean Bronze Age Trade', in N. H. Gale (ed), *Bronze Age Trade in the Mediterranean. Papers Presented at a Conference Held at Rewley House, Oxford, December 1989*. Jonsered: SIMA vol. 90, 351–86.
- Shryock, A., and Smail, D. L. (2011). *Deep History: The Architecture of Past and Present*. Berkeley, CA: University of California Press.
- Sillar, B., and Tite, M. S. (2000). 'The Challenge of "Technological Choices" for Materials Science Approaches in Archaeology', *Archaeometry*, 42:2–20.
- Sindbæk, S. M. (2013). 'Broken Links and Black Boxes: Material Affiliations and Contextual Network Synthesis in the Viking World', in C. Knappett (ed), *Network Analysis in Archaeology: New Approaches to Regional Interaction*. Oxford: Oxford University Press, 71–94.
- Smith, M. E., Feinman, G. M., Drennan, R. D., Earle, T., and Morris, I. (2012). 'Archaeology as a Social Science', *Proceedings of the National Academy of Sciences*, 109(20):7617–21.
- Stein, G. (2002). 'From Passive Periphery to Active Agents: Emerging Perspectives in the Archaeology of Interregional Interaction', *American Anthropologist* 104:903–16.
- Swenson, E. (2012). 'Moche Ceremonial Architecture as Thirdspace: The Politics of Place-Making in the Ancient Andes', *Journal of Social Archaeology*, 12:3–28.
- Tallis, R. (2011). *Aping Mankind: Neuromania, Darwinitis and the Misrepresentation of Humanity*. London: Acumen.
- Van der Leeuw, S. E., and Redman, C. L. (2002). 'Placing Archaeology at the Center of Socio-Natural Studies', *American Antiquity*, 67:597–606.
- Van Oyen, A., (2016). 'Networks or Work-Nets? Actor-Network Theory and Multiple Social Topologies in the Production of Roman *Terra Sigillata*', in T. Brughmans, A. Collar, and F. Coward (eds) *The Connected Past: Challenges to Network Studies in Archaeology and History*. Oxford: Oxford University Press, 35–56.
- Watts, D. J. (2004). 'The "New" Science of Networks', *Annual Review of Sociology*, 30:243–70.
- Watts, D. J., and Strogatz, S. H. (1998). 'Collective Dynamics of "Small-World" Networks', *Nature*, 393:440–2.
- Wellman, B. (1979). 'The Community Question: The Intimate Networks of East Yorkers', *The American Journal of Sociology*, 84:1201–31.

3

Networks or Work-nets? Actor-Network Theory and Multiple Social Topologies in the Production of Roman *Terra Sigillata*

Astrid Van Oyen

INTRODUCTION

This paper has a threefold aim: to clarify some misunderstandings concerning so-called ‘actor–network theory’ (ANT), to show how ANT can be put to use in archaeology, and to articulate differences from, overlaps, and possible combinations with ‘conventional’ network analysis. The key argument is that whereas networks imply direct and untransformed flows between bounded entities, ANT renders visible the heterogeneous ‘work-nets’ needed to support and stabilize these networks and the entities they connect. Moreover, networks are only one constellation or social topology that can emerge from such work-nets, and we thus need to be cautious not to project properties of networks onto other constellations. Finally, ANT introduces a much-needed dynamic approach in which stability and uniformity of categories and flows is the outcome of contingent processes of maintenance work, and not the starting point.

This has important consequences for how we conceive of material culture, and hence of our data, in archaeology: archaeological network analysis tends to assume stable, invariable, and comparable things, traits, or properties. One extreme example of this tendency is so-called *terra sigillata*, a type of Roman pottery that is so easily recognizable and so well-studied that its stability and invariability across contexts and practices is, to a large extent, taken for granted. The problem is that the starting position of *sigillata* as a category already posits a particular, non-neutral social topology, with certain possibilities for action. By analysing changes and continuities in the practices by which this type of pottery was produced in a single site in central France, however, this paper will explore some of the problems with this assumption, and suggest ways of resolving these issues.

NETWORKS AS AN ANTI-CATEGORICAL MOVE

Networks. What's in a word? Following the recent buzz around networks in everyday and academic vocabulary, the answer seems to veer from 'everything' to 'nothing' and back. Stretched from a lightweight intuition to deep-seated theoretical and methodological axioms, the term risks losing all real potential. Cynics are quick to warn that any discipline, theory, or method flying the banner of networks is amalgamated into a melting pot in which implicit, incompatible, or conflicting principles are stirred. This fate threatens to befall so-called ANT, whose very name has been critiqued and, briefly, changed to 'sociology of translation' by its proponents for that very reason (Latour 1999b; 2005:9). Indeed, the label ANT could easily be misread as a theory linking actors as bounded (and often rational and intentional) individuals through social relations making up a larger structure, a network. Whereas this description could qualify as a blueprint for 'classical'¹ network analysis, the quote marks suggest that this is exactly what ANT is *not* about. Nevertheless, I take this confusion as an opportunity for closer engagement, rather than as a warning sign to halt any cross-fertilization between such conflicting principles. However, engagement can only happen on the condition that we critically dissect the differences between classical network analysis and ANT.

Paradoxically, the most convenient way of taking up this exercise is by emphasizing a shared reaction at the heart of both strands, which we could call *anti-categorical*. Classical (social) network analysis in its various guises reacted against an atomistic perspective on agents and their behaviour (Freeman 2004; Newman 2010). It questioned the assumption that categorical variables (e.g. class, gender) alone could account for human action (Emirbayer and Goodwin 1994: 1414). For example, Granovetter's (1985) seminal work in economic sociology rejected the atomism implied by neoclassical economics, and instead studied economic actors as linked into networks of personal relationships (Krippner and Alvarez 2007). Whereas these relationships had been filtered out as inconsequential in the neoclassical economic paradigm, they became constitutive of the possibilities for economic action in the logic of network analysis. The anti-categorical solution of classical network analysis then is to link previously defined categories (entities composed of categorical variables, e.g. people with a particular job, gender,

¹ Coming from the ANT tradition, I admit to having only a partial and simplified grasp of other forms of (social) network analysis (itself a 'collective term' (Brughmans 2010: 278)), which I will therefore label as 'classical'. Moreover, I need to point out that the straw man of classical network analysis set up in this paper draws on problems associated with applications to date, but which are not inherent to network analysis as such. Despite this rather coarse-grained understanding, I hope to be able to draw out some useful differences and potential additions, on which others more versed in the 'classical' tradition can build.

nationality) together. Put differently, connectedness is seen as another variable to be mapped *in addition* to traits such as gender, class (Knox et al. 2006: 118; Malkin et al. 2007: 4).

Similarly, ANT arose out of dissatisfaction with existing explanatory strategies in social theory, which often put forward categorical variables (again gender, class, and so on) as constituting the explanation of action, instead of being in need of explanation themselves (e.g. why do we have a category ‘gender’?) (Latour 2005). But instead of *adding* connections to those other attributes, ANT posits relationality as a way of completely clearing our analytical toolbox of those a priori tags. It does not deny the existence of such categories and variables—as we will explore below—but it urges the analyst to observe their eventual existence take shape as a result of the analysis, instead of presuming it at the start. The network in ANT is thus a metaphor for a general starting principle of relationality. This relational principle in turn serves as a heuristic device for an analysis that is not bound by the limits of predefined variables such as gender or class, or by predefined domains such as science, society, or nature (Latour 1993, 2012). It follows that ANT is inherently different from classical network approaches in that it lets go of any prior assumption of networks as constellations of *stable* entities (e.g. people of certain gender and class; or pots of certain form, as e.g. in Sindbæk 2007) connected by *variable* social ties. Because ANT not only questions the connections between various categorical attributes, but also challenges their very make-up (e.g. how is gender defined?), it can be said to be concerned with ontology (Mol 2002).

FROM NETWORKS TO MULTIPLE TOPOLOGIES

Perhaps the work most often cited as heralding the development of ANT is Latour and Woolgar’s (1979) study of laboratory practices. Having observed the day-to-day events in a scientific laboratory, they argued that the existence of scientific facts (generally said to be the natural ‘out there’ correlates unveiled by science) actually depends on a heterogeneous *network* including such non-scientific things as gossip, test tubes, fundraising, machines, patents, pride. Hence the network of ANT is more accurately described as a *work-net* (after Latour 2005: 129–133): charting all the actors involved that do the work to allow experiments to turn into scientific facts and to enable these facts to travel as stable entities from a specific laboratory to other laboratories, handbooks, and papers. In another seminal study, Law (1986) traced the work-net of fifteenth and sixteenth century Portuguese territorial expansion that relied on ships being able to travel across the seas both as stable, physical entities (the ships had to arrive safely) and semantic entities (the ship and its equipment

had to carry the same meanings, associations, and programme of action once arrived at their destination as they had at their departure in Portugal). This involved ‘documents, devices and drilled people’ (Law 1986: 134), in particular to translate the complex findings of astronomy to workable tools on board the ships that allowed sailors to determine their position in the absence of landmarks. As a consequence, a network for ANT is about the effort and transformations involved in constructing and stabilizing things (facts, ships) and transporting them unchanged. This contrasts with the immediacy, transparency, and direct access associated with networks in both common parlance and classical network analysis, from the railway to the internet (Latour 1999b, 2012: 103 and *passim*).

However, the problem with studies in the early days of ANT is that they have been read as positing the network as a unique or universal social topology (for critiques, see Bloomfield and Vurdubakis 1999; Law 1999; Lee and Brown 1994; McLean and Hassard 2004; Mol and Law 1994). The argument in these early works tended to be that even though the stability of a thing (again: facts, ships) was not *a priori* given, this stability nevertheless proved effective once its work-net was in place, because its wide entanglement made it too costly to alter. This effectively re-inserted the categories from which ANT sought to get away, but as constructed by a work-net rather than as natural givens. Despite this important qualification—often glossed over by critics—these early studies did not present constellations other than networks to which work-nets could give rise. In reaction to this, various ‘second generation’² ANT studies have tried to maintain relationality as an heuristic device that can help to describe various social topologies, of which networks linking categories is but one possibility (Cussins 1996; de Laet and Mol 2000; Law and Mol 2001; Law and Singleton 2005; Mol and Law 1994; Murdoch 1998: 359. See Van Oyen 2013a for an archaeological example).

Let us now turn to archaeology. What difference does this make to how we deal with things in the past? Arguably ‘doing networks in archaeology’ entails ‘doing networks of things’ (Knappett 2005, 2011, 2016 (this volume); Hodder 2012). The prerequisite is then to be able to articulate the relations between things and between humans and things. From ANT’s principle of relationality introduced above it follows that we cannot maintain a predefined, stable essence of things, to which variable meanings could be added by enrolment in different networks. In the current paper I will explore this position by

² A second generation of studies, loosely associated with ANT, continued the concern with relationality and material practices initiated by ANT, but steered away from the rigid overtones of networks, by focusing more broadly on alternative ontologies. This is a temporal rather than an authorial evolution, as later work by the founding fathers of ANT (contributions in Law and Hassard 1999) can also be qualified as ‘second generation’ in this sense (Gad and Bruun Jensen 2010).

focusing on one such thing: *terra sigillata*, a widespread type of Roman table ware, produced and consumed across the temporal and geographical span of the western part of the empire (Bémont and Jacob 1986; Brulet et al. 2010; Poblome et al. 2004). Every Roman archaeologist is familiar with *sigillata*: bright red, shiny pots with highly recognizable standardized forms. Pottery specialists might add the attributes of calcareous clays and sintered slips obtained in an oxidizing firing atmosphere (Picon 2002b). Because these traits are so well-defined, *sigillata* is implicitly assumed to be a stable category, fully defined by this limited package of related but isolatable traits (e.g. the traits of calcareous clays and red colour influence each other, but are nevertheless separate in empirical observation). Moreover, it is taken to be a clearly bounded group, to which a random pot can either belong or not.

This definition would be taken as a neutral starting point (the essence of *sigillata*) by classical network approaches, which would then caution that it does not completely cover the social position (meaning (but see in this volume Rivers 2016)) of these pots. What would be missing, according to such a hypothetical classical network analysis of *sigillata*, are the various social relations linking these pots, together with other (equally previously defined) pots, artefacts, or people. This is the anti-categorical move as implemented by classical network analysis. But as Knox, Savage and Harvey (2006: 134) note with acuity: 'The network thus seems to provide a challenge to rigidity but in use has the effect of reintroducing new kinds of rigidities, in a different form.' The problem is that the starting position of *sigillata* as a category already posits a particular, non-neutral social topology, with certain possibilities for action (Van Oyen 2015). This begs the question of what kinds of past realities this social topology allows us to describe, and, more importantly, what it leaves out (Van Oyen 2012b). Although such partiality is not necessarily wrong, and has proven to be a useful exploratory tool, this article attempts another strategy. Instead of starting from the given *terra sigillata* as a stable category and inserting it in networks of various kinds (e.g. based on consumption, formal resemblance, etc.), an ANT approach reverses the order of the exercise by tracing relational work-nets that then shape a particular kind (social topology) of *terra sigillata*, that could or could not be a category.

This paper seeks to illustrate how this approach could work for archaeology, and how it could help in qualifying some of the shortcomings of classical network analysis. It thereby builds on previous work by Knappett (2002, 2005, 2008, 2009, 2011) at the interstices between things, ANT, and networks, but aims to focus more closely on the 'second generation' ANT lineage and its archaeological leverage. The examples discussed below draw on case studies explored in more detail in the context of my PhD dissertation (Van Oyen 2013b, 2015; forthcoming).

HOW TO ‘DO’ ANT IN ARCHAEOLOGY

Whereas the mixed bag of classical network analysis is united by the use of formal methods, methodology is arguably underdeveloped in the loose ANT tradition. In part a symptom of ANT’s different origin (in science, technology, and society studies, see Law 2008) and of its current platform in material culture studies which is characterized by cherry-picked examples rather than extended case studies (Hicks and Beaudry 2010; Tilley et al. 2006), this lack of consideration of methods is nevertheless in need of revision (Knappett 2005, 2008, 2011, and 2016 (this volume)). The pioneering ANT studies discussed above built their analyses on a largely un-reflexive combination of historical research (Law 1986), technical reports (Latour 1999a; Law 2002), and ethnography as a mix of observation, interviews, and participation (Latour 1999a; Latour and Woolgar 1979). While second generation ANT studies have given due credit to the performativity of research practices (e.g. de Laet and Mol 2000; Law 2004; Mol 2002; Thompson 2005), ironically they have remained silent on the actual methods of data collection used.

Archaeology is often portrayed as a discipline with specific data constraints (e.g. ‘partial’ and ‘mute’ evidence). But archaeologists are actually very adept at tracing connections in their data (e.g. Knappett 2005), and we would do well to scrutinize our rich discipline-specific toolbox in light of ANT’s findings on human–thing relations and their emergence and stabilization. Put differently: can we mobilize established methods such as contextual analysis, statistics, attribute analysis as methods for charting relations of data similarity and difference (cf. Hodder 2012)? Future work should also explore whether this is a possible way of en chaining the formal tools of classical network analysis with the theoretical premises of ANT.

This chapter will explore the compatibility between the relational principles of ANT and one such method now firmly embedded in archaeology: *chaîne opératoire* research (Dobres 2000; Dobres and Hoffman 1994; Edmonds 1990; Schlanger 1994). These studies conceive of technology not as a mechanistic linear development (critique by Pfaffenberger 1992), but as contextual and embodied ways of doing or practices (following Mauss 1950). A loosely interpreted *chaîne opératoire* approach can chart relations of similarity and difference in the production sequence of *terra sigillata*, by analysing the variability in several technological choices (Lemonnier 1993; Van Oyen 2012a). The analytical focus of these technological choices has to be on the contingent nature of specific actions rather than on generic activity labels. For instance, the activity of kiln firing is too generic to contribute much, whereas the contrast between firing at a constant temperature of 1050°C in an oxidizing kiln atmosphere and firing at variable temperatures between 800 and 1000°C in a reducing kiln atmosphere speaks to the variability of practices. Charting this range of variation in various technological choices (e.g. firing

temperatures, clays used), it is possible to gain insight in the social topology of *sigillata* as enacted through production at a certain place and time: how *sigillata* was contingently defined and understood.

Of course any methodology entails artificially pinning down entities and processes—e.g. ‘firing’ itself brings together fuel collection or creation, woodland management, skills, etc. Nevertheless, the reality-producing role thus granted to practice in both method (*chaîne opératoire*) and theory (ANT) (Mol 2002; Wenger 1998) distances this approach from classical network analysis, whose structures often remain ‘abstracted from daily practice and comprehension’ (Knox et al. 2006: 119; but see Graham 2009). A *chaîne opératoire* approach thus seems like a good candidate to be paired with ANT, and to explore how the resulting patterns differ from those to be expected from (a hypothetical) classical network analysis.

EARLY ‘SIGILLATA’ PRODUCTION AT LEZOUX: FROM NETWORK TO WORK-NET

The site of Lezoux (Central Gaul, near present-day Clermont-Ferrand) was the main *sigillata* production centre in the Western Roman Empire in the second century AD. Active as a pottery production site from at least the early first century AD onwards, it offers a long-term perspective able to capture changes in the way *sigillata* was being defined, and thus multiple social topologies. This alone makes it an interesting case study for this research. But what is more, the taken-for-granted dictionary definition of *sigillata* already discussed is an historical construct in large part built on evidence from Lezoux. Seminal works pinning down the key technical parameters of *sigillata* were based on Lezoux products (Picon 1973). With reference to Lezoux, Picon (1973, 2002a) developed a model for *sigillata* production as an opportunistic economic choice: either a cheaper version was produced (fired in a reducing atmosphere), or more investment to produce ‘real’ *sigillata* (fired in an oxidizing atmosphere) paid off because of the possibility of long-distance export. What matters for our purpose is that this model labels these products as two different catalogue entries: clearly bounded, well-defined by a limited package of traits (e.g. firing mode, type of clay), and comparable one to another (the two options can be weighted). To put it very bluntly, all classical network analysis would have to do is to ‘add the social’ by inserting these categories in their respective, variable networks.

However, as we have seen, to ‘add the social’ to artefact classes or categories whose topology is defined prior to the analysis runs against the grain of ANT. Can we instead examine empirically whether the allegedly cheaper *sigillata* version at Lezoux was cast as a strategic, well-defined alternative? The earliest pottery activity at Lezoux itself is attested by a late Iron Age kiln of

the widespread La Tène tradition (Mennessier-Jouannet 1991). Associated ceramics consisted of jar and bowl forms in greyish-black fabrics, and attest to a specialist mastery of modelling on the fast wheel and controlling the firing atmosphere. From around AD 10 onwards (Bet and Delor 2002: 241) evidence of ceramic production at Lezoux increased considerably. By then Italian *sigillata* (Ettlinger et al. 1990) had become a consistent element of imports in the region, and some of those forms were incorporated in the ceramic repertoire of the micaceous (referring to the fabric) Lezoux wares (AD 10–early second century AD); or the cheaper *sigillata*.

Classical network analysis could chart lines indicating stylistic influences, and probably even reflecting the movement of potters (Bet et al. 1994: 43–44; Hartley 1977; Hoffmann and Juranek 1982; Vertet 1967: 257–262) between Lezoux and contemporary *sigillata* production sites in Italy and South Gaul. But such an approach would freeze individual instances of shared potters' stamps or forms and extrapolate from these to a shared topology of *sigillata* as a well-defined category, whose traits (those stamps and forms) could be selected, related, and compared. The crux of the matter is that the resulting network would assume direct access, transparency, and immediacy of flow, without 'paying' for this by highlighting the supporting work-net. In this case, the work-net is about how those influences were locally mediated, made sense of, and defined. Let us look at how this happened.

The site of Lezoux was organized around different production groups, and recent research has shown that this fragmented production landscape existed from the start, rather than being the result of an increase in production capacity during the peak production period of the site (Chuniaud 2002: 247). Domestic, cultic, and funerary evidence (Bet 1988) suggests that the different workshop groups were relatively autonomous communities of practice (Wenger 1998). The various groups produced a wide range of ceramics from the early first century AD onwards, with varying forms and surface finishing (Bet et al. 1994; Bet and Delor 2002; Brulet et al. 2010: 108–111). The forms inspired by *sigillata* imports entered into a local repertoire, as illustrated by an example of the South Gaulish '*sigillata*' form Dragendorff 29 executed at Lezoux with a 'non-*sigillata*' lead-glazed technique (Vertet 1968: 30). Moreover, whatever mechanisms channelled the transmission and influence of forms, it is clear that selective appropriation and transformation were practiced. For example, production of the South Gaulish form Ritterling 9 is not attested at Lezoux (Bet and Delor 2002: 240), whereas a specific variant of the South Gaulish form Dragendorff 29 was developed here (Vertet 1967: 279–283; and 1968). These *transformative* relations of the work-net not only affected the shapes of pots, but also branched into other technological solutions: the Lezoux version of Dragendorff 29 was made using a mould without a base, and the final product's inner surface was never slipped—all in contrast to Italian and South Gaulish practices.

The non-calcareous clays used to produce the first *sigillata* forms at Lezoux did not differ in any way from those employed for other local ceramic products, be they coarse or fine wares (Picon et al. 1971: 195). This continuity was maintained despite the presence of both non-calcareous and calcareous clay beds in the vicinity of Lezoux (Van Oyen 2012a: 53–55). Apparently the choice of the ‘right’ clay was (probably tacitly) key to the embedded pottery tradition to the extent of not allowing for change (what van der Leeuw (1993: 240) calls ‘invariant “backbones”’). This contrasts with the shifting forms and appearances of locally produced ceramics. As a result a clear *difference* from South Gaulish and Italian *sigillata* production was created, which invariably used calcareous clays (e.g. Picon 2002b); and an equally clear *similarity* was maintained with other Lezoux products (or forms). Indeed, a shared use of clay allowed various production sequences to converge throughout a considerable part of the *chaîne opératoire*.

The early *sigillata* forms produced at Lezoux were fired in a reducing atmosphere, followed by an oxidizing cooling phase. This allowed output of an orange-red exterior colour, echoing *sigillata* imports. However, the lack of an oxidizing firing atmosphere in combination with widely diverging maximum temperatures (Brulet et al. 2010: 108; Picon 1973), meant that an even red colour could not be guaranteed, and the slip rarely attained a sintered and shiny aspect as with imported *sigillata*. The circular kilns used in this period derived from the late Iron Age kiln-type of the La Tène tradition already mentioned (Bet et al. 1994: 47). In terms of the work-net, it is likely that this long tradition of know-how combining clay types, kiln technology, and firing mode was too ‘costly’ to change, not in the formalist economic sense implied by Picon (2002a) with alternative options (the catalogue entries) being weighted, but because alternative options were simply not possible, conceivable, or visible within the established understanding of ceramic production. Instead, the latitude of variation within the anchored knowledge system (forms, colours, etc.) was exploited to accommodate the appearance of *sigillata* imports.

It is easy to imagine how a hypothetical classical network analysis, modelled on the applications of network analysis to date in archaeology and history, would take for granted properties of early Lezoux ‘*sigillata*’—in particular its formal resemblance to contemporary Italian and South Gaulish *sigillata*—and use these as a starting point for drawing networks of transfer, influence, and migration. This would not only reify these properties and their possible variation (e.g. it is hard to pigeonhole the local adaptation that we recognize as deriving from the South Gaulish ‘*sigillata*’ form Dragendorff 29) but also the underlying definition of *sigillata* as a bounded category, a package that can be identified based on the presence of some of its traits. Moreover, it would implicitly set up a binary boundary between the *sigillata*-network and the ‘other’ products made at Lezoux. The above analysis, however, has shown that at least in terms of production practices, the enrolment of *sigillata* forms did

not introduce conceptual change: the locally embedded routines of ceramic production went largely unquestioned, and possible links resulting from, for example, the incorporation of ‘South Gaulish’ forms, were not defined or understood as ‘extra-local’ connections.

Let me clarify this with a modern example: the telephone (for a similar case treated from the perspective of management strategy, see Lane and Maxfield 1997). The underlying ontology of classical network analysis would assume a stable essence of ‘telephone’ directly linking the wired phone to the mobile phone, based for instance on shared capacities of use. Following that logic, both products could then be compared, and the mobile phone could be characterized, for example, as ‘a phone without a wire’; much like early Lezoux *sigillata* would be ‘*sigillata* with non-calcareous clays’. If considered from the *work-net* deployed in the production practices of these respective products, however, the initial relation between the mobile ‘phone’ and the wired ‘phone’ has equally little ontological leverage as that between Italian ‘*sigillata*’ and early Lezoux ‘*sigillata*’: while the mobile phone was an offshoot of radio technology and its possibilities, the wired phone hinged on established principles of electrics, and was already embedded in a wide range of consumer practices.

Critics will argue, however, that if we nowadays range the mobile and wired phone under the same heading of ‘phone’, then they must share some sort of essence. The following section will explain how to accept the first part of that deduction, but not the second.

SIGILLATA AS A CATEGORY: FROM ANALYTICAL A PRIORI TO SOCIAL TOPOLOGY

The unchallenged, embodied tradition of ‘how to make good pottery’ (which happened to include forms scholars now label as ‘*sigillata*’) at Lezoux encountered another orthodoxy of ‘how to make good *sigillata* forms’ around the second half of the first century AD. Moreover, the products of this new, incoming orthodoxy looked very similar to the earlier Lezoux *sigillata*. We can only speculate as to how this similarity must have triggered considerable interest, positive or negative. Although initially drawn in by this veneer of similarity in appearance (forms, colour), soon, the differences between the respective practices would have become apparent (Van Oyen 2012a).

The most distinctive difference concerns a switch to calcareous clays, and as a consequence a complete differentiation of the *sigillata* production sequence from other ceramic products, which continued to use non-calcareous clays (Brulet et al. 2010: 324–326; Picon 1973; Picon et al. 1971). Fetching and preparing clay for *sigillata* now became an isolated event in the taskscape (Ingold 2000) of pottery production. This was linked to firing of *sigillata* forms

in an oxidizing instead of a reducing atmosphere, which again created a difference between *sigillata* and other pots in practice—they could no longer be fired in a single kiln load. By the middle of the second century, larger rectangular kilns (Chuniaud 2002: 245; Delage 1998: 286) and extended paved areas for clay preparation testify to changes in technical and organizational practices. These paved areas again highlight the construction of a practical and conceptual boundary between the different types of clay destined for different products and requiring different treatment in a different setting.

But the observed technical transition did not equate to the wholesale adoption of a ‘catalogue entry’ *sigillata* package. Processes of trial-and-error occurred, and earlier practices of producing *sigillata* (non-calcareous clays, reducing firing) persisted at Lezoux until c. AD 140. Again, we must trace a work-net (creating and stabilizing categories) before being able to introduce the possibility of a network as used in current classical network analysis (linking those categories). Here the multifocal production landscape of Lezoux gives some clues. The technical transitions went hand in hand with a spatial relocation of the core of the *sigillata* production activity (Delage 1998: 281). This was paired with the introduction of new potters’ names, representing workshops (e.g. Libertus, Butrio), and possibly also with changing patterns of land ownership, as indicated by replacement of agricultural structures with potting infrastructure in those newly instated workshop groups (Bet 1988).

Around the beginning of the first century, and associated with some of the names of prominent new potters, a wide range of forms and many misfired ceramics point to a period of experimentation. Experimentation is also evident in the search for new techniques to obtain new effects, such as plaster moulds leaving a dotted imprint on the vessel surface, which were, however, rapidly abandoned (Bet et al. 1987: xii). This buzz of trial-and-error testifies to an almost analytical preoccupation with the relationship between techniques, bodily actions, and end product. It follows that choices were not random, but were made in a discursive way. Where empirical links were not necessarily understood as potentially meaningful connections in the case of early Lezoux *sigillata* described in the previous section, this changed when faced with the new ways of producing *sigillata*. Previously invisible boundaries of the imaginable and the possible were made visible, questioned, and stretched. A similar process played out in the relationship between those same new potters’ names and the short-lived phenomenon of black *sigillata* (Simpson 1957). These technologically equivalent but differently fired pots comprised a range of *sigillata* forms—barring the most common forms that were increasingly at the heart of what constituted a *sigillata* pot (Symonds 1992: 10)—and some non-*sigillata* profiles such as liquid containers (Brulet et al. 1999: 126). Hence similarity and difference between products and production sequences were open for negotiation.

Around the middle of the second century, the mass of *sigillata* production shifted back to the original production group, where the output of certain workshops increased to a never previously seen level of production (Delage 1998). From that time onwards, the variety of forms shrank notably, with a number of forms becoming especially popular against the backdrop of a standardized repertoire (Bet and Delor 2002; Brulet et al. 2010: 112 ff.). As to decoration, too, greater efficiency and repetition of certain models and canons characterized the second century. Larger workshops with more continuous decorative schemes replaced the rapid turnover of the output of the previous smaller workshops (Brulet et al. 2010: 118–119). Maximum firing temperatures similarly became less variable (Picon 1973).

This reduced latitude of variation made *sigillata* into a *category* at Lezoux in the course of the second half of the second century AD. It could now be fully defined by a limited number of traits; any sherd or pot *either* belonged to it *or* fell outside of its boundaries (note the difference with the issue of analytical boundaries discussed by Peeples et al. 2016). *Sigillata* was now a well-defined and clearly understood kind of thing: something from which you could produce x number of bowls and sell y number of plates without having to specify over and over again how they should be made, which clays should be used, or what they should look like. It also enabled comparison between different specimens within this category and between their respective traits (Van Oyen 2015). In other words, it created the conditions of comparability that are often assumed in classical network analysis. But as Mol (2002: 70) warns: ‘Coordination into singularity doesn’t depend on the possibility to refer to a preexisting object. It is a task.’ Comparability, boundedness, definiteness, etc. are not characteristics or attributes of a world ‘out there’, but are to be made, with a lot of effort (work-net) and develop only in some cases (Law 2004; Latour 2012). In those cases, one can speak of categories, which could in turn be inserted in the networks of classical network analysis. But we can no longer subscribe to such categories as neutral, analytical, and *a priori* now that we have seen that they are a non-neutral kind of social topology, creating certain possibilities for action (e.g. comparability) while precluding others.

The example of the wired telephone and the mobile illustrates a similar process of category formation. Despite their origin in different technological traditions, both products have aligned around a shared niche of private and professional communication. In practice, mobile phones have become an extension of the traditional ‘hotspots’ of home and work where wired telephones provide an anchored communication address. A single operator can now provide packages catering for both mobile and fixed connectivity. This alignment, which is at once practical (e.g. shared operators, bills, address books) and conceptual (both devices become part of the same concepts of connectivity and communication), creates conditions of comparability, both between mobile phones and wired phones (e.g. a mobile phone then effectively

becomes a ‘telephone’ without a wire, amenable to different actions, functions, identities), and within the category of mobile phones itself. For example, it enables comparison based on individual traits; e.g. does the phone come with a camera or not? 3G or 4G? Internet connection or not? Finally, this possibility of comparing traits as elements of a predefined, stable package could then be the basis for a classical network analysis, to inquire for instance whether these categorical variables are in any way related to the interconnectedness between these devices or their users (e.g. Do friends tend to use the same operator? Is there a correlation between gender and the use of a camera on a phone?).

BOUNDARIES AND BOUNDARY WORK

Once categories are recognized as a particular kind of social topology instead of a pre-existing analytical structure, their persistence is no longer guaranteed. Put differently, not only is a category constructed and stabilized by a work-net, the ‘lifelines’ spun by this work-net have to be enacted and re-enacted. This is where the critique by the second generation of ANT has corrected the claims that early ANT studies made all too easily: that once categories (e.g. scientific facts, Portuguese ships) were constructed, they were too widely entangled to be destabilized. Instead, boundary work is needed to maintain the parameters of a category as social topology: a clear boundary demarcating belonging from non-belonging (cf. Bloomfield and Vurdubakis 1999) and a limited package of traits. (Hodder 2012 lays similar stress on how humans are trapped into maintaining things in the particular way they define them.) Maintenance of these parameters does not however rule out empirical changes, for example, in the contents of the respective traits (e.g. forms), as long as the boundaries are maintained in practice.

This begs the question of whether and how the category boundary of *sigillata* was maintained at Lezoux. The workshops experimenting with *sigillata* in the first half of the second century (e.g. those associated with Libertus) went on to produce another type of pottery around the middle of that century (Bet and Gras 1999; Desbat and Vilvorder 2000: 177). These have been rather unhelpfully called ‘Rhenish’ wares (Symonds 1992), and, unlike *sigillata*, their identifying traits are much debated among specialists (e.g. Brulet et al. 1999). The argument of this chapter is that this lack of scholarly conclusion about Rhenish wares is a direct consequence of the social topology of these things. More specifically, Lezoux Rhenish wares were not so much defined positively—i.e. as an internally coherent and externally bounded category—but negatively: as what *sigillata* was not, as *sigillata*’s ‘Other’.

On the one hand, the respective production sequences of Lezoux *sigillata* and Rhenish wares overlapped significantly. They were probably produced by

the same workshops and potters, and deposits of Rhenish wares at the production site are always associated with more numerous fragments of *sigillata* (Bet and Gras 1999: 35). But a more emphatic overlap was the shared use of the same type of clays (Bocquet 1999: 216; Brulet et al. 2010: 345–346). That this was not a purely pragmatic choice prompted by increased efficiency or cost reduction is evidenced by the larger variation in calcium oxide composition in Rhenish ware fabrics, showing that even though the same clays were used, they were treated differently for the separate products (Picon 1973; Bocquet 1999: 223–225). The resulting pattern is one of overlap between both production sequences, followed by a marked divergence. These ‘othering’ mechanisms served to reinforce the boundedness of *sigillata*—and thus its social topology as a category.

Similar ‘othering’ mechanisms can be traced for other technological choices too. For example, both Lezoux *sigillata* and Rhenish wares are counted among tablewares, but the latter tapped into an established non-*sigillata* repertoire of drinking forms. Nevertheless, they also took on some forms derived from the by now standardized *sigillata* repertoire (Bet and Gras 1999: 26–31; Brulet et al. 2010: 346–347; Desbat and Vilvorder 2000: 178). As for decorative techniques, Rhenish wares tended to be given sub-slip trails of barbotine (Bet and Gras 1999: 33–34; Brulet et al. 2010: 346–347; Symonds 1992: 17–26). This necessitated not only a special treatment of the clays, but also specialized skills, different to those of *sigillata* mould makers. Moreover, the viscous nature of barbotine (Symonds 1992: 13) introduced a distinct range of possible decorations, more fluid but less plastic and detailed, with a dominance of the vegetal motives that had by then disappeared from the *sigillata* repertoire. Nevertheless, the occasional use of ‘*sigillata*’ moulding on Lezoux Rhenish wares emphasizes that the barbotine decoration of Rhenish wares was different *in comparison to sigillata*.

This boundary was most sharply defined through the differing modalities of firing and their results. Rhenish wares were fired in a reducing atmosphere (Bocquet 1999: 223–225) and thus could not conceivably be fired in the same batch as *sigillata*. Moreover this resulted in a black exterior colour, in stark contrast to the red surface of *sigillata*. This recalls the black *sigillata* discussed above. Here it becomes clear that the empirical variability on which classical network analysis tends to be based is not coextensive with meaningful patterning. This same instance of artefact variability (black/red) has a very different resonance whether it is articulated within a topology of experimentation (black *sigillata*), or as a boundary marker (black Rhenish wares as different from red *sigillata*). Importantly, this changed leverage is not due to any inherent, essential difference between Rhenish wares and black *sigillata*, but due to the changed social topology of *sigillata*, which had become a bounded category.

Finally, the fact that the boundary work served to establish and confirm *sigillata* as a category and not Rhenish ware is further illustrated by the attestation of forms exclusive to Rhenish wares but with a red exterior colour (see comment by Philippe Bet in Brulet et al. 1999: 125), similar to the examples of moulded Rhenish wares. This is in turn underlined by significant latitude of variation among Rhenish ware production: for example, although always markedly different from *sigillata*'s bright red, surface colour could vary from dark black through brown to green (Brulet et al. 2010: 346; Symonds 1992: 18). This returns to the scholarly issues encountered with pinning down criteria for identifying Rhenish wares, which did not exist as a category but as an 'Other' to a category: defining the boundary around a package of traits, but itself not defined as a bounded package of traits.

Back once more to the telephone example. On the one hand, the work-net that has managed to make the mobile phone into a category has entailed such a wide entanglement that it becomes difficult to challenge this social topology: mobile phones have become deeply embedded in and are constitutive of contemporary business practices, friendships, revolutions, etc. But it does not follow that the mobile phone is now essentialized and its stability guaranteed. On the contrary, what a mobile phone is (the contents of the package of traits) and how it can be defined (as a package of traits/category or another social topology) is constantly challenged, negotiated, and renegotiated in relation to smartphones, tablets, etc.; to the point where 'telephone without a wire' hardly describes a bounded, clearly-defined category anymore.

CONCLUSIONS

Although it is difficult to do justice to the nuances of both ANT and classical network analysis (admittedly a crude banner) in the space permitted, this chapter has drawn out some differences between both strands. This was done not with the aim of polarizing debates, but from the position that a clear articulation of goals, assumptions, and shortcomings is crucial to any attempt at combination or cross-fertilization.

The key difference articulated here is that between a 'work-net' of transformations that keep a certain entity in place on the one hand and a 'network' assuming direct, transparent, and untransformed connection between previously defined entities on the other. Whereas a work-net is built through *practice*, a network is often seen as *impacting* on practices rather than *resulting from them*. For example, networks and their structures have been successfully invoked to account for the spread of religious ideas (Collar 2007, 2011), diseases, or other kinds of 'novel ideas or beliefs' (see in this volume Kandler

and Caccioli 2016), adding to the inherent properties of the things spreading (from religious belief to *sigillata*). But these network analyses do not discuss how the relational configuration of the thing that is being circulated both interacts with and changes how that thing is being defined. The complexities of such a picture, where changing definitions and altered conditions of possibility in turn feed back into how the thing under study can spread, are revealed only when moving beyond the split between the inherent properties and the relational (or network) position of things.

It is increasingly recognized in classical network analysis that network measures (e.g. centrality) and structural positions do not determine action but set negotiable conditions of possibility (see in this volume Düring 2016). Similarly, according to second generation ANT, different social topologies entail different possibilities for action. Although not elaborated in this chapter, the eventual definition and maintenance of *sigillata* as a category impacted on its further trajectory by shaping the possibilities for reproduction, distribution, and consumption (Van Oyen 2015). These possibilities were distinct from those triggered by other social topologies as for instance early Lezoux ceramic production or Lezoux ‘Rhenish’ wares. In a similar vein, we have seen how the same empirical relation of ‘black’ to ‘red’ went from an instance of experimentation as part of a wider process of negotiating topologies to a boundary marker, as a result of the transformation of *sigillata* into a category. Do beliefs, diseases, or opinions similarly change shape and properties while travelling through the networks upheld by the practices of heterogeneous work-nets? How did this shape-changing affect these practices, and in turn the networks they stabilized?

The problem with classical network analysis from an ANT perspective is not just that the possible interpretive claims are already (in part) contained in the analytical toolbox (from a network analysis perspective, see Butts 2009)—which is the case for any methodology—but that what is presupposed is precisely what is the subject of inquiry for ANT. It is as if network analysis skips the steps that are key to ANT: tracing the work-net that enacts a certain social topology. As a result, classical network analysis runs the risk of missing out on a wide range of possible social topologies that are not compatible with its method and with the very template of the network. For example, it arguably would have mistaken the empirical ties of stylistic influence in early Lezoux *sigillata* for meaningful extra-local references, when in fact it seems that these were very much part of an unchallenged local way of making sense of ceramic production. If anything, ANT can thus caution classical network analysis about the level of taken-for-grantedness in its data which, if not critically scrutinized, threatens to give rise to generic stories where ‘*sigillata* is *sigillata* is *sigillata*’, based on generic networks where *sigillata* in place *x* is by definition connected with *sigillata* in place *y* (e.g. Sindbæk 2007). I would thus tentatively suggest that it would be helpful to take into account an ANT-approach as

developed in this chapter when defining the categories and components of classical network analysis (Brughmans 2010: 298 and *passim*).

This brings us back to the problem of building networks based on things, or, how to do networks in archaeology. If human agents are put at the forefront as both the building blocks of the networks and the means of tracing those networks analytically, then active verbs can often qualify the relations or edges: A ‘knows’, ‘speaks to’, ‘befriends’, or ‘cites’ B. But what are the correlates of these verbs when dealing with *things* (others have pondered a similar issue: Jones 2001; Knappett 2008)? Here ANT is particularly promising for archaeology, as it denies a divide between (passive) objects and (active) subjects (Latour 1993, 2000, 2005). Instead, it provides us with a nuanced model of how things work, emphasizing material practices (Mol 2002) and defining things *in-the-doing*. As shown in this chapter, these principles prove compatible with some well-established disciplinary tools in archaeology, such as *chaîne opératoire* research. In contrast, ‘classical’ network analysis as it currently stands threatens to reify the process of the creation and stabilization of categories in which ANT is interested. Whether network analysis could be modified as a tool for exploring the concerns of ANT is a matter for further investigation by scholars better versed in its complexities.

ACKNOWLEDGEMENTS

I am indebted to Martin Millett, Carl Knappett, John Robb, and the editors for reading and commenting on previous drafts. Any remaining errors are my own.

REFERENCES

- Bémont, C., and Jacob, J.-P. (eds) (1986). *La terre sigillée gallo-romaine. Lieux de production du Haut-Empire. Implantations, produits, relations* (*Documents d'Archéologie française* 6). Paris: Editions de la Maison des Sciences de l'Homme.
- Bet, P. (1988). ‘Groupes d'ateliers et potiers de Lezoux (Puy-de-Dôme) durant la période gallo-romaine’, *SFECAg Actes du congrès d'Orange*: 221–41.
- Bet, P., Delage, R., and Vernhet, A. (1994). ‘Lezoux et Millau. Confrontation d'idées et de données’, *SFECAg Actes du congrès de Millau*: 43–61.
- Bet, P., and Delor, A. (2002). ‘Les premiers ateliers céramiques de type méditerranéen en Auvergne, l'exemple des officines de sigillée’, in L. Rivet, and M. Sciallano (eds), *Vivre, produire et échanger. Reflets méditerranéens. Mélanges offerts à Bernard Liou* (*Archéologie et histoire romaine* 8). Montagnac: Monique Mergoil, 235–42.
- Bet, P., Gangloff, R., and Vertet, H. (1987). *Les productions céramiques antiques de Lezoux et de la Gaule centrale à travers les collections du Musée archéologique de Lezoux* (Revue archéologique SITES. Hors-série 32). Gonfaron: AFAM.

- Bet, P., and Gras, D. (1999). 'Parois fines engobées et céramique métallescente de Lezoux', in R. Brulet, R. P. Symonds, and F. Vilvorder (eds), *Céramiques engobées et métallescentes gallo-romaines. Actes du colloque organisé à Louvain-la-Neuve le 18 mars 1995 (Rei Cretariae Romanae Fautorum Acta Supplementum 8)*. Oxford: RCRF, 13–38.
- Bloomfield, B. P., and Vurdubakis, T. (1999). 'The Outer Limits. Monsters, Actor Networks and the Writing of Displacement', *Organization*, 6:625–47.
- Bocquet, A. (1999). 'La production et la distribution des céramiques fines engobées et métallescentes dans le nord de la Gaule. Approche minéralogique et géochimique', in R. Brulet, R. P. Symonds, and F. Vilvorder (eds), *Céramiques engobées et métallescentes gallo-romaines. Actes du colloque organisé à Louvain-la-Neuve le 18 mars 1995 (Rei Cretariae Romanae Fautorum Acta Supplementum 8)*. Oxford: RCRF, 129–286.
- Brughmans, T. (2010). 'Connecting the Dots. Towards Archaeological Network Analysis', *Oxford Journal of Archaeology*, 29:277–303.
- Brulet, R., Symonds, R. P., and Vilvorder, F. (eds) (1999). *Céramiques engobées et métallescentes gallo-romaines. Actes du colloque organisé à Louvain-la-Neuve le 18 mars 1995 (Rei Cretariae Romanae Fautorum Acta Supplementum 8)*. Oxford: RCRF.
- Brulet, R., Vilvorder, F., and Delage, R. (eds) (2010). *La céramique romaine en Gaule du Nord. Dictionnaire des céramiques. La vaisselle à large diffusion*. Turnhout: Brepols.
- Butts, C. T. (2009) 'Revisiting the Foundations of Network Analysis', *Science*, 325: 414–16.
- Chuniaud, K. (2002). 'Le groupe des ateliers de potiers de Ligonnes à Lezoux (Puy-de-Dôme). Un champ d'étude pour les questions relatives à l'organisation de la production céramique en Gaule romaine', in L. Rivet and M. Sciallano (eds), *Vivre, produire et échanger. Reflets méditerranéens. Mélanges offerts à Bernard Liou (Archéologie et histoire romaine 8)*. Montagnac: Monique Mergoil, 243–48.
- Collar, A. C. F. (2007). 'Network Theory and Religious Innovation', *Mediterranean Historical Review*, 22:149–62.
- Collar, A. C. F. (2011). 'Military Networks and the Cult of Jupiter Dolichenus', in E. Winter (ed.), *Von Kummuh nach Telouch. Historische und archäologische Untersuchungen in Kommagene. Dolichener und Kommagenerische Forschungen IV (Asia Minor Studien, 64)*. Bonn: Habelt, 217–47.
- Cussins, C. (1996). 'Ontological Choreography. Agency through Objectification in Infertility Clinics', *Social Studies of Science*, 26:575–610.
- de Laet, M., and Mol, A. (2000). 'The Zimbabwe Bush Pump. Mechanics of a Fluid Technology', *Social Studies of Science*, 30:225–63.
- Delage, R. (1998). 'Première approche de la diffusion des céramiques sigillées du centre de la Gaule en occident romain', *SFECAg Actes du congrès d'Istres*: 271–313.
- Desbat, R., and Vilvorder, F. (2000). 'Die Trierer Schwarzfurnisware. Produktion und Handel', in K. Strobel (ed.), *Forschungen zur römischen Keramikindustrie. Produktions-, Rechts- und Distributionstrukturen. Akten des 1. Trierer Symposiums zur antiken Wirtschaftsgeschichte (Trierer historische Forschungen, 42)*. Mainz: von Zabern, 177–91.

- Dobres, M.-A. (2000). *Technology and Social Agency. Outlining a Practice Framework for Archaeology*. Oxford: Blackwell.
- Dobres, M.-A., and Hoffman, C. R. (1994). 'Social Agency and the Dynamics of Prehistoric Technology', *Journal of Archaeological Method and Theory*, 1:211–58.
- Düring, M. (2016). 'How Reliable are Centrality Measures for Data Collected from Fragmentary and Heterogeneous Historical Sources? A Case Study', in Brughmans T., Collar A., and Coward F. (eds), *The Connected Past: Challenges to Network Studies in Archaeology and History*. Oxford: Oxford University Press, 85–101.
- Edmonds, M. (1990). 'Description, Understanding and the *chaîne opératoire*', *Archaeological Review from Cambridge*, 9: 55–70.
- Emirbayer, M., and Goodwin, J. (1994). 'Network Analysis, Culture, and the Problem of Agency', *American Journal of Sociology*, 99:1411–54.
- Ettlinger, E., et al. (eds) (1990) *Conspectus Formarum Terrae Sigillatae Italico Modo Confectae*. Bonn: Habelt.
- Freeman, L. C. (2004). *The Development of Social Network Analysis*. Vancouver: Empirical Press.
- Gad, C., and Bruun Jensen, C. (2010). 'On the Consequences of Post-ANT', *Science, Technology and Human Values*, 35:55–80.
- Graham, S. (2009). 'The Space Between. The Geography of Social Networks in the Tiber Valley', in F. Coarelli and H. Patterson (eds), *Mercator Placidissimus. The Tiber Valley in Antiquity. New Research in the Upper and Middle River Valley. Proceedings of the Conference held at the British School at Rome, 27–28 February 2004*. Rome: Edizioni Quasar, 671–86.
- Granovetter, M. (1985). 'Economic Action and Social Structure. The Problem of Embeddedness', *American Journal of Sociology*, 91:481–510.
- Hartley, B. R. (1977). 'Some wandering potters', in J. Dore and K. Greene (eds), *Roman Pottery Studies in Britain and Beyond. Papers Presented to John Gillam, July 1977 (BAR Supplementary Series 30)*. Oxford: British Archaeological Reports, 251–61.
- Hicks, D., and Beaudry, M. C. (2010). 'Introduction', in D. Hicks and M. C. Beaudry (eds), *The Oxford Handbook of Material Culture Studies*. Oxford: Oxford University Press, 1–21.
- Hodder, I. (2012). *Entangled. An Archaeology of the Relationship between Humans and Things*. Malden (MA): Wiley-Blackwell.
- Hoffmann, B., and Juranek, H. (1982). 'Bestätigung der Zusammenhänge von La Graufesenque und Lezoux durch chemische und töpferische/technische Analyse des Abdrucks eines Bildstempels', *Rei Cretariae Romanae Fautorum Acta*, 21–22:7987.
- Ingold, T. (2000). *The Perception of the Environment. Essays in Livelihood, Dwelling and Skill*. London: Routledge.
- Jones, A. (2001). 'Drawn from memory. The archaeology of aesthetics and the aesthetics of archaeology in earlier Bronze Age Britain and the present', *World Archaeology*, 33:334–56.
- Kandler, A., and Caccioli, F. (2016). 'Networks, Homophily and the Spread of Innovations', in T. Brughmans, A. Collar and F. Coward (eds), *The Connected Past: Challenges to Network Studies in Archaeology and History*. Oxford: Oxford University Press, 175–97.

- Knappett, C. (2002). 'Photographs, Skeuomorphs and Marionettes. Some Thoughts on Mind, Agency and Object', *Journal of Material Culture*, 7:97–117.
- Knappett, C. (2005). *Thinking through Material Culture. An Interdisciplinary Perspective*. Philadelphia (PA): University of Pennsylvania Press.
- Knappett, C. (2008). 'The Neglected Networks of Material Agency. Artefacts, Pictures and Texts', in C. Knappett and L. Malafouris (eds), *Material Agency. Towards a Non-Anthropocentric Approach*. New-York: Springer, 139–56.
- Knappett, C. (2009). 'Meaning in miniature. Semiotic networks in material culture', in M. Jessen, N. Johannsen and H. J. Jensen (eds), *Excavating the Mind*. Aarhus: Aarhus University Press, 87–109.
- Knappett, C. (2011). *An Archaeology of Interaction. Network Perspectives on Material Culture and Society*. Oxford: Oxford University Press.
- Knappett, C. (2016). 'Networks in Archaeology: Between Scientific Method and Humanistic Metaphor', in T. Brughmans, A. Collar, and F. Coward (eds), *The Connected Past: Challenges to Network Studies in Archaeology and History*. Oxford: Oxford University Press, 21–33.
- Knox, H., Savage, M., and Harvey, P. (2006). 'Social networks and the study of relations. Networks as method, metaphor and form', *Economy and Society*, 35:113–40.
- Krippner, G. R., and Alvarez, A. S. (2007). 'Embeddedness and the Intellectual Projects of Economic Sociology', *Annual Review of Sociology*, 33:219–40.
- Lane, D., and Maxfield, R. (1997). 'Foresight, Complexity, and Strategy', in W. B. Arthur, S. N. Durlauf, and D. A. Lane (eds), *The Economy as an Evolving Complex System II*. Reading (MA): Perseus Books, 169–98.
- Latour, B. (1993). *We Have Never Been Modern*. Cambridge (MA): Harvard University Press.
- Latour, B. (1999a). *Pandora's Hope. Essays on the Reality of Science Studies*. Cambridge (MA): Harvard University Press.
- Latour, B. (1999b). 'On recalling ANT', in J. Law and J. Hassard (eds), *Actor Network Theory and After*. Oxford: Blackwell, 15–25.
- Latour, B. (2000). 'The Berlin key or how to do words with things', in P. M. Graves-Brown (ed.), *Matter, Materiality, and Modern Culture*. London: Routledge, 10–21.
- Latour, B. (2005). *Reassembling the Social. An Introduction to Actor-Network-Theory*. Oxford: Oxford University Press.
- Latour, B. (2012). *Enquête sur les modes d'existence. Une anthropologie des Modernes*. Paris: La Découverte.
- Latour, B., and Woolgar, S. (1979). *Laboratory Life. The Construction of Scientific Facts*. Beverly Hills-London: Sage.
- Law, J. (1986). 'On the Methods of Long Distance Control. Vessels, Navigation and the Portuguese Route to India', in J. Law (ed.), *Power, Action and Belief. A New Sociology of Knowledge? (Sociological Review Monograph 32)*. London: Routledge and Kegan Paul, 234–63.
- Law, J. (1999). 'After ANT. Complexity, naming and topology', in J. Law and J. Hassard (eds), *Actor Network Theory and After*. Oxford: Blackwell, 1–14.
- Law, J. (2002). *Aircraft Stories. Decentering the Object in Technoscience*. Durham (NC) and London: Duke University Press.

- Law, J. (2004). *After Method. Mess in Social Science Research*. London: Routledge.
- Law, J. (2008). 'On sociology and STS', *The Sociological Review*, 56(4):623–49.
- Law, J., and Hassard, J. (eds) (1999). *Actor Network Theory and After*. Oxford: Blackwell.
- Law, J., and Mol, A. (2001). 'Situating Technoscience. An Inquiry into Spatialities', *Environment and Planning D. Society and Space*, 19:609–21.
- Law, J. and Singleton, V. (2005). 'Object Lessons', *Organization*, 12(3):331–55.
- Lee, N., and Brown, S. (1994). 'Otherness and the Actor Network. The Undiscovered Continent', *American Behavioral Scientist*, 37:772–90.
- Lemonnier, P. (1993). 'Introduction', in P. Lemonnier (ed.), *Technological Choices. Transformation in material cultures since the Neolithic*. London: Routledge, 1–35.
- Malkin, I., Constantakopoulou, C., and Panagopoulou, K. (2007). 'Preface. Networks in the Ancient Mediterranean', *Mediterranean Historical Review*, 22:1–9.
- Mauss, M. (1950). *Sociologie et anthropologie*. Paris: Presses universitaires de France.
- McLean, C., and Hassard, J. (2004). 'Symmetrical absence/symmetrical absurdity. Critical notes on the production of actor-network accounts', *Journal of Management Studies*, 41:493–519.
- Mennessier-Jouannet, C. (1991). 'Un four de potier de La Tène Finale à Lezoux (Puyde-Dôme)', *Revue archéologique du Centre de la France*, 30:113–26.
- Mol, A. (2002). *The Body Multiple. Ontology in Medical Practice*. Durham (NC) and London: Duke University Press.
- Mol, A., and Law, J. (1994). 'Regions, Networks and Fluids. Anaemia and Social Topology', *Social Studies of Science*, 24:641–71.
- Murdoch, J. (1998). 'The Spaces of Actor-Network Theory', *Geoforum*, 29:357–74.
- Newman, M. E. J. (2010). *Networks. An Introduction*. Oxford: Oxford University Press.
- Peeples, M. A., Mills, B. J., Haas, W. R. J., Clark, J. J., and Roberts, J. M. J. (2016). 'Analytical Challenges for the Application of Social Network Analysis in Archaeology', in T. Brughmans, A. Collar, and F. Coward (eds), *The Connected Past: Challenges to Network Studies in Archaeology and History*. Oxford: Oxford University Press, 59–84.
- Pfaffenberger, B. (1992). 'Social Anthropology of Technology', *Annual Review of Anthropology*, 21:491–516.
- Picon, M. (1973). *Introduction à l'étude technique des céramiques sigillées de Lezoux (Université de Dijon. Faculté des sciences humaines, Centre des recherches sur techniques gréco-romaines 2)*. Dijon: Presses de l'imprimerie universitaire.
- Picon, M. (2002a). 'A propos des sigillées, présigillées et imitations des sigillées. Questions de "coûts" et de marchés', *SFECAG, Actes du congrès de Bayeux*: 345–56.
- Picon, M. (2002b). 'Les modes de cuisson, les pâtes et les vernis de la Graufesenque. Une mise au point', in M. Genin and A. Vernhet (eds), *Céramiques de la Graufesenque et autres productions d'époque romaine. Nouvelles recherches. Hommages à Bettina Hoffmann (Archéologie et Histoire romaine 7)*. Montagnac: Monique Mer-goil, 139–63.
- Picon, M., Vichy, M., and Meille, E. (1971). 'Composition of the Lezoux, Lyon and Arezzo Samian Ware', *Archaeometry*, 13:191–208.
- Poblome, J., Talloen, P., Brulet, R., and Waelkens, M. (eds) (2004). *Early Italian Sigillata. The Chronological Framework and Trade Patterns. Proceedings of the First International ROCT-Congress Leuven, May 7 and 8, 1999*. Leuven: Peeters.

- Rivers, R. (2016). 'Can Archaeological Models Always Fulfil Our Prejudices?', in T. Brughmans, A. Collar, and F. Coward (eds), *The Connected Past: Challenges to Network Studies in Archaeology and History*. Oxford: Oxford University Press, 123–47.
- Schlanger, N. (1994). 'Mindful technology. unleashing the *chaîne opératoire* for an archaeology of mind', in C. Renfrew and E. B. W. Zubrow (eds), *The Ancient Mind. Elements of Cognitive Archaeology*. Cambridge: Cambridge University Press, 143–51.
- Simpson, G. (1957). 'Metallic black slip vessels from Central Gaul with applied and moulded decoration', *The Antiquaries Journal*, 37:29–42.
- Sindbæk, S. M. (2007). 'Networks and nodal points. The emergence of towns in early Viking Age Scandinavia', *Antiquity*, 81:119–32.
- Symonds, R. P. (1992). *Rhenish Wares. Fine Dark Coloured Pottery From Gaul and Germany* (Oxford University Committee for Archaeology 23). Oxford: Oxford University Committee for Archaeology.
- Thompson, C. (2005). *Making Parents. The Ontological Choreography of Reproductive Technologies*. Cambridge (MA) and London: MIT Press.
- Tilley, C., Keane, W., Küchler, S., Rowlands, M., and Spyer, P. (eds) (2006). *Handbook of Material Culture*. London: Sage.
- van der Leeuw, S. (1993). 'Giving the potter a choice. Conceptual aspects of pottery techniques', in P. Lemonnier (ed.), *Technological Choices. Transformation in material cultures since the Neolithic*. London: Routledge, 238–88.
- Van Oyen, A. (2012a). 'Knowledge systems in the production of terra sigillata. Moving beyond the local/global paradox', in M. Duggan, F. McIntosh, and D. J. Rohl (eds) *TRAC 2011. Proceedings of the Twenty First Theoretical Roman Archaeology Conference, Newcastle 2011*. Oxford: Oxbow, 48–59.
- Van Oyen, A. (2012b). Review of C. Knappett (2011) *An Archaeology of Interaction. Network Perspectives on Material Culture and Society*, Oxford: Oxford University Press, *Archaeological Review from Cambridge* 27(2), 220–28.
- Van Oyen, A. (2013a). 'Towards a postcolonial artefact analysis', *Archaeological Dialogues*, 20:79–105.
- Van Oyen, A. (2013b). 'Rethinking terra sigillata. An archaeological application of Actor-Network Theory', Ph.D. thesis (University of Cambridge).
- Van Oyen, A. (2015). 'Actor-Network Theory's Take on Archaeological Types: Becoming, Material Agency, and Historical Explanation', *Cambridge Archaeological Journal*, 25(1), 63–78.
- Van Oyen, A. (forthcoming). *How Things Make History. The Roman Empire and its Terra Sigillata Pottery*. Amsterdam: Amsterdam University Press.
- Vertet, H. (1967). 'Céramique sigillée tibérienne à Lezoux', *Revue archéologique*, 2:255–86.
- Vertet, H. (1968). 'Influence des céramiques italiennes sur les ateliers arvernes au début du I^{er} siècle', *Revue archéologique du centre*, 7:25–34.
- Wenger, E. (1998). *Communities of Practice. Learning, Meaning and Identity*. Cambridge: Cambridge University Press.

Part II

Challenging Network Analysis of Archaeological and Historical Data

4

Analytical Challenges for the Application of Social Network Analysis in Archaeology

*Matthew A. Peeples, Barbara J. Mills, W. Randall Haas, Jr,
Jeffery J. Clark, and John M. Roberts, Jr*

INTRODUCTION

As the chapters and citations in this volume attest, applications of network analytical techniques using archaeological data have a great deal of potential for both addressing traditional archaeological questions and for providing new directions for archaeological research. Importantly, however, many of the network models and methods imported from other fields that are currently gaining popularity within archaeology (see Brughmans 2013) have not yet been fully assessed in relation to the unique strengths and constraints of archaeological data. We argue that archaeological applications of network analyses necessitate particularly careful consideration of the nature of the data included and the applicability of network metrics, many of which were designed with quite different time-scales and levels of certainty in mind. We further argue that, if archaeologists are able to overcome such challenges, the opportunities afforded by archaeological data (e.g. long-term perspectives, material perspectives) will allow us to contribute substantially to broader interdisciplinary discussions of network methodology, interpretation, and theory.

In this chapter, we explore four general challenges facing archaeologists applying formal network methods: (1) the use of artefacts to construct network relations; (2) temporal variation among the units of analysis; (3) the definition of network boundaries; and (4) the impact of incomplete datasets. We have confronted each of these issues in our own archaeological network analyses focused on prehispanic settlements in the US Southwest. Some challenges are unique to archaeology while others have been extensively discussed in other disciplinary contexts. Following a brief overview of the data and methods that form the basis of our study, we summarize each of these

major issues and offer suggestions for how they might be addressed based on our own experiences and analyses. We do not suggest that we cover all or even most of the challenges archaeological network analysts will face. We do, however, offer a framework for assessing other factors that may influence characterizations of archaeological networks and a general approach for tempering interpretations in relation to such factors.

THE SOUTHWEST SOCIAL NETWORKS PROJECT

The case study that forms the basis of the examples we use in this chapter comes from the Southwest Social Networks (SWSN) Project. The SWSN Project is an interdisciplinary, collaborative effort involving specialists in archaeology, sociology, geographic information systems (GIS), geochemistry, and computer science. Using network methods and models along with geographic and artefact data, our team is exploring the topology and dynamics of social networks in the US Southwest in relation to other archaeologically documented transformations. We have focused thus far on the late prehispanic period (AD 1200–1500), which was a period of dramatic demographic and social change. The SWSN database consists of geographic, site-size, chronological, and architectural information for over 1700 major settlements (> 12 rooms) in Arizona and New Mexico west of the North American Continental Divide (see Fig. 4.1; these data comprise a subset of the larger Coalescent Communities Database; see Hill et al. 2004, 2012; Wilcox et al. 2003). We have also compiled systematic tabulations of painted and plain ceramic types and wares for over 700 of these sites and sourced obsidian objects for nearly 150 sites from published sources, unpublished notes, and new analyses conducted by team members. Altogether, the database contains information on more than 4.3 million ceramics classified by type and 6,000 chemically characterized obsidian objects representing the results of over a century of archaeological research in the region.

Creating a Network

Using the settlement and artefact data described above, we have developed a set of related methods for constructing archaeological networks using material culture data. In the examples provided in this chapter, we focus on methods for creating networks from systematic decorated ceramic ware counts. As they are defined in the US Southwest, ceramic wares (which often consist of finer type designations) are large sets of related ceramics, grouped according to similarities in both design and technology. Wares are relatively easy to distinguish, even from small fragments, and are often produced within relatively




Figure 4.1. Map of the Southwest Social Networks Project study area.

limited geographic areas. Our basic argument, which we expand upon further in a subsequent section (Mills et al. 2013a, 2013b), is that similarities in the wares used and discarded at two sites can be used to assess the strength of social connections among those sites, whether direct or indirect, material or informational. Such similarities can be driven by a number of social processes including exchange, emulation, population movement, and especially, shared contexts of production or learning. Assemblage-level similarities in the proportions of ceramic wares result from the combined effects of all of these processes and can provide a general indication of the strongest connections among settlements, especially at regional scales.

To define ceramic networks, we treat every site in the SWSN database with systematic decorated ceramic data as a node. Because our ceramic data come from sites spanning more than three centuries, many with non-overlapping occupations, we cannot directly compare ceramic assemblages until we first account for differences in occupation length and period. Our team has developed a method of chronologically apportioning assemblages into short temporal intervals (each 50-year interval from AD 1200 to 1500 in this case) using the dates associated with ceramic types present at a site and the estimated

occupation span of that site. We describe this procedure in more detail below (Roberts et al. 2012). The result of this ceramic apportioning procedure is a series of tables for each 50-year interval between AD 1200 and 1500, consisting of the estimated frequencies of each ceramic ware deposited during that interval. This apportioning procedure allows us to make comparisons among sites that have dramatically different occupation spans, and also allows us to explore changes in network properties and structures through time.

Using the chronologically apportioned ceramic data for each of our six 50-year intervals, we next create a series of similarity matrices among sites in our database for each interval. For this example, we use the Brainerd–Robinson similarity coefficient (Brainerd 1951; Robinson 1951), where similarity (S) between site a and site b is defined as:

$$S_{ab} = \frac{200 - \sum_k |P_{ak} - P_{bk}|}{200} \quad (\text{Eq. 4.1})$$

k = all ceramic wares; P_{ak} = per cent of ware k at site a ; P_{bk} = per cent of ware k at site b .

This is, of course, only one of numerous similarity/distance metrics we could have selected (e.g. chi-square distance, Bray–Curtis dissimilarity). We selected the Brainerd–Robinson coefficient because it is relatively intuitive, is commonly used for comparing archaeological assemblages, and because it works well for nominal percentage data (e.g. Cowgill 1990; see also Golitko et al. 2012; Hart and Engelbrecht 2012; Mills et al. 2013a; Peeples 2011; Peeples and Haas 2013). Brainerd–Robinson similarity coefficients (referred to hereafter as similarity scores) typically range between 0 and 200, but we rescale this measure here to range from 0, indicating no similarity, to 1, indicating perfect similarity. This change is partially cosmetic but, more practically, it simplifies the calculation of other network properties. The result of this step is a symmetric similarity matrix for each 50-year interval with values between 0 and 1 indicating the relative similarities of each site to every other site in our database.

We interpret the similarity scores calculated above as a proxy for the strength of connections among pairs of sites in our network. In network parlance, these are the weights of ties or edges between nodes (sites). These weighted ties are undirected, meaning that we make no assumptions about the direction of flows between nodes. In some circumstances, it is also useful to create binary (present or absent) ties between sites in our network. There are a number of different ways to define binary ties from weighted relational data (Peeples and Roberts 2013). In the examples presented here, we simply define a tie between two sites as present when those sites have a ceramic similarity score of 0.75 or greater. We see the creation of such binary ties as primarily an aid in visualizing these data rather than an analytical step. Wherever possible,

we quantify network properties, such as centrality scores, using the raw similarity scores rather than this binary representation.

Once these networks have been created, numerous node- and graph-level indices can be calculated to characterize different aspects of network structure (Borgatti 2005; Wasserman and Faust 1994). In the context of this chapter, we use two common measures of centrality: degree and eigenvector centrality. Degree centrality for a node in a binary network is calculated simply as the total number of active ties for that node. Where weighted ties are available, degree centrality can be defined as the sum of the weights for all of a node's ties (Newman 2001; Opsahl et al. 2010). This measure provides a general indication of the importance of a node in directing flows through a network in terms of first-order (direct) connections (Freeman 1978). The second measure we use here, eigenvector centrality, is based on the first eigenvector of the adjacency matrix underlying a network. In the case of weighted ties, a node's centrality score is proportional to the weighted (by the tie strengths) sum of the centrality scores of its neighbours (Mills et al. 2013a). A node will have a higher eigenvector centrality score when it is connected to many other highly central nodes. This measure assumes that a node can influence all other nodes in the network simultaneously rather than simply in terms of direct ties. We have found eigenvector centrality quite useful in our own work (Mills et al. 2013a, 2013b), and others have shown that the measure provides a good characterization of centrality based on the structure of a network as a whole, especially for large information networks such as ours, in which the direction and exact routes of transmission cannot be determined (Bonacich 1972, 1991; Borgatti 2005).

ANALYTICAL CHALLENGES

The archaeological data that support our research on prehispanic social network structure in the US Southwest have required careful consideration in relation to their interpretative potential and the constraints of archaeological networks. Together, these potentialities and constraints present theoretical and practical challenges for archaeological network analysts. In this section, we briefly explore four of these challenges: (1) defining material-based networks; (2) assessing temporal variation; (3) defining network boundaries; and (4) dealing with incomplete data.

Measuring Social Connections through Material Culture

One of the first considerations for any material-based archaeological network study is deciding what artefacts or other attributes should be used to define network ties among the units of analysis. Alternatively, we might ask what our

archaeological data can and cannot tell us about the social networks of interest (see also Knappett 2016). Ultimately, the answer to these questions depends on the specific problem being investigated and the quality of available information. In this section we discuss the advantages and disadvantages of several common forms of artefact/attribute data that have recently been used in the creation of networks in a number of archaeological settings (e.g. source determinations; presence/absence data; frequency data). We suggest that these different kinds of data are variously useful for defining network ties relating to different social processes, in particular material procurement, production, and consumption. In this discussion, we emphasize that it is important to always consider exactly what a given dataset is telling us about connections between our units of analysis and to use analyses that play to the strengths of those data.

In many ways, artefacts that can be sourced to a specific location of origin through compositional analyses or other related methods are ideal candidates for archaeological network analyses. Source determinations allow us to explore the connections between settlements or regions based on the movement of objects or the shared use of raw materials. In contexts where source determinations can be made at relatively fine geographic scales, it may even be possible to actually characterize the volume of flows of materials among sites or regions in different directions. Network terminology, visualizations, and the network metaphors have long been a part of compositional and other provenance studies (e.g. Evans 1989), but the application of formal network analytical methods and characterizations of network properties are a relatively new development (Bernardini 2007; Clark et al. 2012; Golitko et al. 2012).

Through the SWSN project, we have compiled a large body of x-ray fluorescence (XRF) chemical characterization data for thousands of obsidian artefacts from sites across our study area (Clark et al. 2012). Obsidian is well-suited to chemical characterization as individual sources are chemically distinct, easily distinguished through multiple methods, and can often be defined with a great deal of geographic specificity. In the US Southwest and Northern Mexico, XRF characterization has been used to define fifty chemically distinct and geographically known sources of obsidian that are also found in archaeological contexts (Shackley 2005). Our current work explores the relationship between the relative proportions of objects from different obsidian sources found at settlements and the costs associated with travel from the settlements to those obsidian source areas (Clark et al. 2012). Sourced obsidian objects also form the basis of a recent study of Maya obsidian exchange by Golitko and colleagues (2012). In their analysis, similarities in the relative frequencies of obsidian sources found at sites are used to define network ties among those sites. As these examples illustrate, sourced obsidian works well as a basis for defining networks in terms of both two-mode (site to source) and one-mode (site to site) connections. In either case, however, it is important to think about how the objects used to define such ties are procured. For example, do

similarities in the obsidian sources represented at two sites indicate direct exchange between those sites, procurement of raw materials from the same source, down-the-line exchange of materials, or some other social process? Such questions typically need to be addressed with multiple lines of evidence and the answers may vary for different sites or sources within a single network. It is important to keep in mind that ties in such a network, even when source areas are known with reasonable certainty, may not all have the same underlying basis.

Ceramics are another common material class that has been subjected to compositional and provenance studies in the US Southwest and around the world, based on a variety of techniques including petrography, neutron activation analysis, and lead isotope analyses (e.g. Abbott 2000; Glowacki and Neff 2002; Huntley et al. 2012). In contrast to obsidian, the results of these analyses largely focus on production rather than procurement areas. Due to the nature of ceramic technology and composition, source determinations are often much coarser than those of materials like obsidian and, thus, it can be difficult to track the direct movement of vessels between geographically specific source areas and destinations (but see Bernardini 2007; Schachner et al. 2011). Further, unlike obsidian data, it is often difficult to combine results based on different analytical methods because the resolution at which production areas can be defined will often vary considerably. Such data do have a high potential for archaeological network analyses, however, particularly when geographic sources are well-defined and samples are large enough and designed in such a way as to allow for considerations of the volume of ceramic materials moving from production areas to consuming settlements (e.g. Bernardini 2007). At the same time, it is important to consider that similarities in compositional groups present at two sites do not necessarily indicate direct interaction and may variously represent the exchange of vessels, population movement, the use of geologically similar materials, or some combination of these factors. Again, untangling these different sources of variation typically requires multiple lines of material evidence.

Another common approach to creating material-based archaeological networks is to use the presence/absence or frequency of different artefact types as the basis for ties. This approach could entail, for example, creating ties between two sites (or other units of analysis) when they both share some distinct class of object (Mizoguchi 2009) or by weighting ties based on the number of artefact categories that nodes share (Brughmans 2010; Coward 2010). This method for defining ties can be restricted to types or categories within a single material class (e.g. ceramic wares; see Brughmans 2010) or can include multiple material classes simultaneously (Coward 2010). Such simple classifications of ties have the advantage of allowing the use of data from diverse sources, even those where the particular collection strategies are not fully known. Alternatively, where systematic data, including the relative frequencies of different objects or classes of objects are available, it is possible to

define ties and the weights of the ties among nodes in terms of assemblage level similarities measures. It is this approach, described above, which forms the basis of most of the SWSN project analyses. The major advantage of this approach is that it allows for a consideration of both the diversity and the richness of objects or types in the creation of ties between nodes. Indeed, different similarity metrics can be selected to differentially emphasize rare categories (e.g. chi-square distance) or the most common categories (e.g. Brainerd–Robinson similarities). The systematic data required for such an approach may be considerably more difficult to obtain than presence/absence data, however, especially at large scales. Thus, such an approach may limit the sample of sites or other units available for consideration.

In general, the approaches to network creation involving the presence, frequency, or similarity of objects described above are focused more on the consumption (i.e. patterns of use and discard) of objects rather than their production or the procurement of raw materials. In our own work with the SWSN database, we have conceptualized the connections between sites in terms of similarities in the consumption of ceramic wares as indicative of ‘communities of practice’ that are ubiquitous. In the US Southwest, ceramic wares are defined based on technological styles that include the choices of clays, tempers, and surface treatments (e.g. slip colour and pigment types). Southwestern archaeological assemblages, especially after AD 1000, often contain a diversity of wares in different proportions. Similarities in ceramic assemblages are not random, but result from a number of social processes including not only exchange and population movement but also the transmission of knowledge about ceramic production, and other learned behaviours, such as the socially appropriate ways to cook, serve, and store food and other goods. We do not argue that high similarity in discarded ceramics always indicates direct interactions among the residents of two sites. We do, however, suggest that the residents of settlements with similar ware assemblages were more likely to have interacted with each other and with more frequency than the residents of settlements who used and discarded quite different sets of ceramic wares.

Because of differences in materials chosen for analysis and the flows that structure their distribution, the nature of the ties in each of the above cases is variable and different kinds of analyses may be applied. For example, obsidian might be regarded as a ‘package’ that moves through the network, often directly from a source to a specific destination. In the case of the decorated ceramics, however, similarities are based on the flow of information or beliefs and attitudes about what the correct way to make a pot is, how it should be decorated, what it is used for, and how it is consumed, among other processes. These different kinds of flows have been discussed by Borgatti (2005) because they can determine the most appropriate choices of centrality measures. In the first example, obsidian, the objects spread through transfer along relatively

well-defined vectors while in the second, they spread through influence, which is more diffuse, and flows under a ‘walk’ process in which the connections linking two nodes can be circuitous, and individual nodes and edges may be visited multiple times along the way. In the first case, it might be more appropriate to use degree centrality or related measures based on directed ties whereas in the latter case it may also be appropriate to use eigenvector centrality (Borgatti 2005). As these examples illustrate, it is important to consider what processes are likely to have been behind the flows of material or information that underlie a given network and to use the appropriate methods for defining ties and characterizing the structure of that network.

Accounting for Temporal Differences

One challenge that is particularly germane to archaeological and historical applications of network analyses is the need to account for temporal differences among the units of analysis. Sites or other units in a network, under many typical applications and conceptualizations of network flows, can only be directly connected if they are contemporaneous. In contexts characterized by sites with quite long occupation spans, this may be difficult to determine definitively. Further, when connections between nodes in an archaeological network are based on material culture, especially in terms of the relative frequencies of objects, it can often be difficult to determine what portions of the total artefact assemblages at the sites we wish to compare actually overlapped in time.

One common approach that archaeologists have used to deal with contemporaneity is to create networks using existing archaeological phase designations. This approach works well for classes of artefacts that are difficult to date directly. Although such an approach has the advantage of not pushing the data beyond what can be reasonably supported, such coarse resolution limits our interpretations of network dynamics. Further, phases of long duration risk inclusion of sites that were not in fact contemporaneous. This problem is not unique to network analyses, as archaeologists often want to untangle complex chronological relationships in archaeological assemblages, for example, for the purposes of seriation (e.g. Kohler and Blinman 1987).

In our own research, we have developed a method for apportioning (or dividing) ceramic assemblages into intervals of a given length using available chronological information for both sites and ceramic types. Using methods described in detail by Roberts and colleagues (2012), we apportion the total ceramic assemblages from each site into each 50-year interval between AD 1200 and 1500. Briefly, this method of ceramic apportioning uses the estimated occupation span of a site, along with the dated production ranges of the types found at that site, and assumptions regarding the prevalence of types through




Figure 4.2. Overview of the ceramic apportioning procedure. Reused with permission from Roberts et al. 2012.

time (we assume that type frequencies approximate a normal curve through time) to estimate the relative proportions of each ware that were deposited in each 50-year interval in which that site was occupied (see Fig. 4.2). This procedure not only allows us to make comparisons among sites that have different occupation spans, but also allows us to explore changes in network properties and structures through time.

Although the apportioning procedure described above greatly improves our ability to compare sites with overlapping, but not identical, date ranges, the apportioned data sets may be somewhat sensitive to the assumption that a normal curve describes a ware's historical popularity (Roberts et al. 2012), and further are only as good as the information we are able to supply regarding site

occupation and ceramic type date ranges. When either of these two factors is in question, it is good practice to test various scenarios as a check on how our initial assumptions may drive our final interpretations. To illustrate this approach, we explore the effects of estimated date range for one type in our ceramic database (Tanque Verde Red-on-brown). Tanque Verde Red-on-brown, a type within Tucson Basin Brown Ware, is found in many late prehispanic settlements in the southern US Southwest, but is most common in the vicinity of the city of Tucson, Arizona (Heckman 2000: Figure 41). Although most archaeologists generally agree that the production of Tanque Verde Red-on-brown began at about AD 1150, opinions vary widely regarding the end date as there are few unambiguous absolute dates available. Based on available dates and associations with better dated types, some archaeologists argue that Tanque Verde Red-on-brown production ceased around AD 1350 or possibly even earlier, while others see it extending all the way to the end of the Hohokam sequence in the southern Southwest at around AD 1450 (compare Deaver 1989; Heckman 2000; Martynec 1993; Wallace 1995). As the jury is still out on which end date is most appropriate, we must consider how these possibilities may influence our interpretations of network dynamics. In order to test the differing assumptions regarding the production dates, we apportioned our ceramic data into 50-year intervals assuming two possible end dates for Tanque Verde Red-on-brown (AD 1350 and AD 1450). We then created similarity matrices and binarized networks using the procedures described above for each interval using both potential end dates for sake of comparison.

Figure 4.3 shows network diagrams for the southern half of the Southwest, including the area around Tucson, for the two fourteenth-century 50-year intervals with both estimated end dates for Tanque Verde Red-on-brown. Although the binarized networks are shown here, the individual nodes are scaled based on their eigenvector centrality scores from the full weighted network. As these plots show, the networks for the AD 1300–1350 interval are quite similar for both the AD 1350 and AD 1450 end dates. The Tucson basin and nearby areas, which were the centres of production for Tucson Basin Brown Ware, form a distinct component with few strong similarities to other sites in the southern Southwest. Sites in the Tucson area almost all have low eigenvector centrality scores, suggesting that they may be somewhat peripheral to interactions characterizing sites across the region as a whole. During the AD 1350–1400 interval, however, we see substantial differences in the networks based on different end dates for Tanque Verde Red-on-Brown. If we assume that the type continued to be produced until AD 1450, the Tucson area sites continue to show low eigenvector centrality and remain relatively distinct from other sites in the southern Southwest. In contrast, if Tanque Verde production ceased at AD 1350, the few remaining sites in the Tucson area are highly central and well integrated into the large network component that characterizes most of the sites in the southern Southwest. These two




Figure 4.3. Network diagrams for the southern Southwest for the AD 1300–1350 and AD 1350–1400 intervals assuming different production dates for Tanque Verde Red-on-brown. Sites located in the Tucson Basin/Papagueria area are shown in black.

different date range scenarios suggest different social processes and relationships for a large portion of our study area.

Overall, the brief example presented in this section suggests that contemporaneity as well as our assumptions regarding ceramic and site dates are important and can dramatically influence network characteristics and our interpretations of network structure. The question of the end dates associated with Tanque Verde Red-on-brown will need to be addressed through additional archaeological investigation and cannot be resolved directly through network analyses. By testing various assumptions with our network data, however, we are able to assess the degree to which any results we obtain are a function of our initial assumptions. Such an exercise is important as it allows us to explore various possibilities while not putting too much weight on results that may be simply a product of a single, and potentially inaccurate, inference.

Defining Network Boundaries

Defining the boundaries of enquiry is an important step in every archaeological investigation. Early archaeological projects were largely site based, and many remain that way. Regional analyses have recently become much more

common although the spatial scale of the regions being investigated and how they are defined vary widely. In many cases, archaeologists do not have a choice in bounding their project areas because they are working within the limits of projects that are predefined by other analysts or governed by heritage management considerations. The issue of boundary definition raises a number of related pragmatic questions for network analyses. What spatial and social scales are most appropriate for archaeological network studies (e.g. intra-site, regions, river valleys, continents)? Do archaeological networks require spatial boundaries or could some other criteria be used? What are the effects of varying network boundaries? These are all issues of boundary specification which have also been much discussed by sociologists in the application of social network analysis in other settings (Laumann et al. 1992).

Sociologists distinguish between networks made up of nodes that share commonalities based on easily recognized, substantial, and well-defined relationships such as spatial relationships or familial ties versus those networks that are defined by more ephemeral criteria such as shared membership in associations or externally (researcher) defined categories. Laumann and others (1992: 70) refer to these different kinds of networks as ‘realist’ versus ‘nominal’ respectively. In archaeology, spatially defined (realist) networks might encompass a set of architectural features within a single site (e.g. Pailes 2014) or spatially clustered sites, or communities, or river valleys and other physiographic units. Nominal networks can be bounded based on attributes of the sites (nodes) in question themselves. For example, archaeologists may use the presence of a particular kind of public architectural feature at a site as the basis for bounding a network, and then examine relationships among all of the sites with that feature independently through other material remains. Both the realist and nominalist approaches are defensible in light of network models and methods as networks, unlike social groups, ‘have no “natural” boundaries’ (Borgatti and Halgin 2011: 1169).

Archaeologists frequently discuss ‘multiscalar’ approaches to their research and this is no different when talking about networks (Knappett 2016). We have taken a multiscalar approach in the SWSN project and have found that at each scale our questions and our interpretations shift (Mills et al. 2015). At what we describe as the ‘micro-scale,’ we have examined how social networks changed at the scale of a single river valley or drainage basin. The San Pedro River valley in southeastern Arizona is one example that we have studied extensively through network analysis (Mills et al. 2013a). Here there are well-documented settlements occupied by the ‘first comers’ or hosts and others that were built somewhat later by migrants who moved from north-eastern Arizona in the late thirteenth century (Clerk and Lyons 2012; Haury 1958; Neuzil 2008; Woodson 1999). We explored how networks changed before, during, and after migration and found that migrants dramatically altered local network topology. Host settlements that had the highest centrality (eigenvector in




Figure 4.4. Network diagram of the greater Hohokam region for the AD 1250–1300, 1300–1350, and 1350–1400 intervals. Sites with public architectural features are shown as squares and sites without these features are circles. Points are scaled based on their eigenvector centrality scores.

this case) early on were those located in the best-watered locations. Many of these highly central sites were also those that persisted longer. Through time, however, migrant settlements became increasingly central to flows measured within this network, at least in part because they produced a widely distributed ceramic ware consumed at settlements throughout the valley (Salado polychrome; see Mills et al. 2013a, 2015).

At what we refer to as the ‘meso-scale’ we consider connections among sites across much of the broader southern Southwest (Fig. 4.4). This network incorporates the San Pedro valley, but also several other valleys and basins lying south of the Mogollon Rim (a major physiographic boundary) in southern Arizona and New Mexico. At this scale we can see how, during the AD 1250–1300 interval, each river valley largely comprised a separate network component. After AD 1300, network components in this region began to coalesce almost entirely into a single large component linked by shared consumption of the same decorated ceramics. The largest connected component in this network after AD 1300 is largely based on high frequencies of Salado polychrome pottery. Salado polychrome vessels are decorated with designs that had strong ideological content (Crown 1994). The large size of the bowls of this ware and the presence of decoration on the highly visible exterior suggest that these vessels were used at least in part in feasting (Mills 2007). Their widespread distribution, the social contexts of their use, and their ideological content has been attributed to the growing popularity of a religion (Crown 1994), or a social movement. At the meso-scale, we are able to see how the increasing connectivity within the San Pedro Valley was actually just a small part of a broader trend toward the creation of strong connections among sites throughout the southern Southwest, driven in large part by the emergence and spread of Salado polychrome.




Figure 4.5. Network diagrams of the entire Southwest Social Networks Project study area for the AD 1250–1300, 1300–1350, 1350–1400, and 1400–1450 intervals. Nodes for sites located in the southern Southwest are shown in black and those in the northern Southwest are shown in white. Points are scaled based on their eigenvector centrality scores.

We refer to the largest scale we have examined to date as the ‘macro-scale’ (Fig. 4.5; see Mills et al. 2013b, 2015). This scale, which constitutes our entire current study area, includes settlements above and below the Mogollon Rim throughout the western half of the Southwest. When we consider both the northern and southern Southwest together, differences in the characteristics and trajectories of the networks across the study area are readily apparent. First, with the exception of a small number of sites in areas lying along the boundary between the north and the south, these two areas appear to be characterized by few strong similarities in ceramic assemblages for any interval. The northern Southwest is characterized by densest connections, the largest connected components, and sites with the highest centrality during the 1250–1300 interval. After AD 1300, however, we see a shift as settlements in the south increase in centrality and the number of ties and size of components increases dramatically, while in the north, dense connections begin to dissolve. Since this is after the period of migration from the northern to the southern Southwest described above, these changes in network topology are likely not coincidental. The

divergence of network trajectories increases through the remaining periods in our study. As this brief example illustrates, at the macro-scale we are able to see dramatic changes in network characteristics in relation to a major period of migration in both the source and destinations of those migrants. Explorations of networks at this macro-scale allow us to relate changes in network structure to larger demographic processes that may not be apparent when focused at smaller scales.

As the brief discussion above highlights, the selection of different network boundaries should be tied to archaeologically significant questions that are appropriate to the social and spatial scales of interest. Different scales of networks can provide different insights into social processes at work in a given setting. The ability to vary network boundaries does not make one scale less valid than another. It does, however, mean that the kinds of questions that can be asked and answered may be quite different.

Assessing the Impact of Partial Datasets

In a perfect world, we would base our interpretations of network properties on the ‘whole network’ including all possible nodes and all ties between those nodes. Although often presented as the ultimate goal of many network data-gathering efforts, complete networks of this sort are almost never attainable at large scales. Even for populations with relatively well-delineated boundaries where high quality information is available, data may be missing for a number of reasons including, for example, the refusal of individuals to participate or the imperfect memories of respondents (Brewer 2000; Granovetter 1976). In other contexts, such as studies of covert networks including terrorist organizations (Carley et al. 2001, 2003) or large and dynamic networks, such as those based on website or email links (Gjoka et al. 2011; Leskovec and Faloutsos 2006), we may actually have very little information on the size and scale of missing data. Most archaeological contexts in which network analyses have been applied probably approximate these more extreme cases of missing information (we do not know what we do not know). Adding to this, in many archaeological contexts our ability to sample a complete network is also often limited by factors outside of our immediate control (e.g. what sites have already been excavated, what collections have been analysed, jurisdictional boundaries). Thus, for large networks in general and archaeological or historical networks in particular, it is important to consider the potential impact of missing information, as any conclusions drawn based on characteristics derived from an incomplete network are subject to some degree of uncertainty (see in this volume Düring 2016; Tsirogiannis and Tsirogiannis 2016).

There is already a large and robust body of literature focused on the effects of sampling and missing information on many different kinds of network data

(Bolland 1988; Borgatti et al. 2006; Costenbader and Valente 2003; Galaskiewicz 1991; Lee et al. 2009; Marsden 1993; Rivera-Hutinel et al. 2012). This past research suggests that different measures of node centrality and graph level network properties respond to missing data in different ways and to widely varying degrees depending on the underlying characteristics of the network in question (see overview in Costenbader and Valente 2003). Although a few overarching patterns emerge in comparisons of network sampling across many contexts (e.g. smaller samples are typically associated with greater variability), there are few ‘rules’ in terms of what measures are appropriate or how much missing data is allowable to accurately characterize a given network. Thus, the approach many researchers have taken is to explore the potential effects of missing information on a case-by-case basis, using the available data as a guide. Although it is impossible to directly measure the effects of missing data if we do not have the full network as a source of comparison, it is possible to assess the stability of various network measures to similar perturbations in the available data using resampling methods (see Mills et al. 2013b, supplement).

In the following exercise, we generally follow methods and heuristics developed in previously published studies of network sampling (Costenbader and Valente 2003; Galaskiewicz 1991) to test the potential influence of missing nodes on patterns we observe in our SWSN ceramic networks. Specifically, we created a series of replicates of each 50-year interval between AD 1200 and 1450 by sub-sampling sites from each network with sampling fractions between 10 per cent and 90 per cent (at 10 per cent intervals) of the total number of sites for that interval. Through this procedure, we created 1000 replicates for each time period and for each sampling fraction for a total of 45000 simulated networks (1000×5 intervals $\times 9$ sampling fractions). We then calculated degree and eigenvector centrality scores for the sampled sites in every simulated network, and for each replicate calculated product-moment correlations between these centrality scores and the scores for these sites in the original data. We also calculated the standard deviation of these correlation coefficients to provide a means for tracking variability at different sampling fractions and time periods.

Figures 4.6 and 4.7 show the mean correlation coefficients for each time period by sampling fraction for degree and eigenvector centrality respectively. As these figures illustrate, correlations between the actual and sub-sampled networks are quite high for both degree and eigenvector centrality in most cases, even when the sampling fraction is as small as 10 per cent of the original. They also show that mean correlations do decrease as the sampling fraction decreases. In all but one set of replicates for either centrality measure, however, the mean correlation between actual node centrality and centrality in the replicates is greater than 0.81. The one exception to this pattern is a correlation of 0.64 for eigenvector centrality for the 10 per cent sampling fraction in the




Figure 4.6. Correlations between degree centrality in full network for each 50-year interval and sub-samples at sampling fractions from 90% to 10% of total.




Figure 4.7. Correlations between eigenvector centrality in full network for each 50-year interval and sub-samples at sampling fractions from 90% to 10% of total.

Table 4.1. Standard deviations of correlations between centrality scores in full network and sub-samples of various sampling fractions.

		Sampling Fraction									
		90%	80%	70%	60%	50%	40%	30%	20%	10%	
Degree Centrality	AD 1200–1250	0.001	0.003	0.005	0.008	0.012	0.017	0.031	0.057	0.138	
	AD 1250–1300	0.000	0.001	0.002	0.003	0.005	0.009	0.015	0.032	0.094	
	AD 1300–1350	0.002	0.004	0.008	0.013	0.021	0.028	0.051	0.087	0.202	
	AD 1350–1400	0.000	0.001	0.002	0.004	0.006	0.012	0.034	0.061	0.194	
	AD 1400–1450	0.000	0.001	0.002	0.003	0.006	0.014	0.036	0.125	0.321	
Eigenvector Centrality	AD 1200–1250	0.000	0.001	0.003	0.006	0.025	0.070	0.125	0.224	0.413	
	AD 1250–1300	0.000	0.000	0.001	0.001	0.002	0.003	0.008	0.037	0.233	
	AD 1300–1350	0.002	0.007	0.019	0.044	0.073	0.119	0.205	0.289	0.440	
	AD 1350–1400	0.000	0.000	0.000	0.001	0.001	0.004	0.011	0.072	0.298	
	AD 1400–1450	0.000	0.000	0.001	0.001	0.002	0.005	0.065	0.148	0.379	

AD 1300–1350 interval. Correlations for AD 1300–1350 tend to decline as the sampling fraction decreases at a somewhat faster rate than other intervals considered for both degree and eigenvector centrality. Interestingly, this interval marks a particularly dynamic period in the demographic and social history of our study area characterized by aggregation and a fundamental reorganization of communities, large-scale population movement, and the establishment of clusters of large settlements separated by vast buffer zones (Hill et al. 2004). This association suggests that the rapid pace of change and the diversity of relationships characterizing different portions of the study area at this time may make the network associated with this dynamic interval particularly susceptible to problems associated with missing data.

Table 4.1 shows the standard deviations of correlation coefficients for each interval and each sampling fraction. This measure provides a general indication of how sampling fraction (i.e. the amount of missing data) might influence the degree of variation in centrality scores. As this table indicates, the standard deviations for both eigenvector and degree centrality are quite low in most cases until the sampling fraction reaches 20 per cent or 10 per cent. In general, standard deviations increase as the sampling fraction decreases. Further, it appears that the total size of the sample (number of sites in full dataset) also matters. As might be expected, time periods with fewer sites tend to show greater variability in correlations of centrality at low sample fractions. The standard deviations of scores increase more rapidly as the sampling fraction decreases for the AD 1300–1350 interval, again suggesting that the rapid change occurring during this period may make it susceptible to missing data.

In general, the results presented here are encouraging. Even when only a small sample of the total network is included, the strongest patterns in relative centrality among sites in our sample are still evident. At the same time, small

differences in centrality among sites do appear to be subject to variation due to the inclusion or exclusion of specific sites, especially when the total number of sites is small and the sampling fraction is small. This suggests that such minor differences for either of these centrality measures should not be interpreted as particularly meaningful. At the same time, these results suggest that moderate to large differences in centrality, especially at the regional scale, are quite robust to missing data. Based on these results and related tests we have conducted in our other analyses, we limit our interpretations of patterns in our networks to those characterized by moderate to large differences among sites and we do not rely heavily on particular values or network metrics for any specific site.

As many datasets used for network analyses in archaeology are samples of a larger universe, we advocate using similar resampling methods to those described above to assess the stability of centrality measures or other network metrics to missing data. Further, similar approaches could also be useful in assessing the effects of other potential confounding factors. For example, we have used a bootstrapping approach to assess the influence of sampling variability in ceramic assemblages on centrality measures (Mills et al. 2013a). Such a procedure can aid considerably in separating real patterning in network data from ‘noise’ introduced by sampling effort and variability and other sources of uncertainty at a variety of scales.

CONCLUSIONS

The discussion above has highlighted just a small number of the many challenges that face researchers interested in applying formal network models and methods to archaeological data. Luckily, these challenges are not insurmountable, but do suggest the need for a careful consideration of research design, the methods to be used, and the assumptions involved in the creation of network datasets. Based on the issues discussed here, we conclude with a few suggestions that will be relevant for many challenges that are likely to be encountered by those engaged in archaeological network studies.

First, the examples described above illustrate that it is always important to keep in mind exactly what constitutes a tie in a given network and what these ties can and cannot tell us about interaction and other social processes. The edges of networks are often glossed simply as ‘connections,’ but the discussion above suggests that edges may indicate quite different kinds of relationships among nodes within a single network or between networks (see also in this volume Van Oyen 2016). For example, the social processes driving the movement of obsidian objects from a source area to a destination may be quite different from the processes that generate similarities in ceramic ware

assemblages. Further, as various network properties are appropriate for different kinds of relations and flow processes (Borgatti 2005), it is important to match the analyses used with the data at hand. In archaeological networks connections are often based on similarities in material proxies with the exact character of the interaction remaining unclear, as opposed to the less ambiguous evidence for direct and directed connections among nodes often available in many contemporary network analyses. Thus, network models and methods that assume network flows are direct or efficient may not be appropriate in many archaeological case studies.

Beyond this, the examples provided here suggest that it is useful to test the sensitivity of any substantive results against potential sources of variation or error in a network dataset. One example above suggests that the strength of connection between sites in the Tucson basin and other areas in the southern Southwest is strongly dependent on our assumptions regarding the production dates for a single type (Tanque Verde Red-on-brown). Until new data are available to resolve the production date for that type, it is appropriate to limit our interpretations of a given network to patterns that do not hinge on this single assumption alone. We advocate such investigation as a way of identifying those graph- and node-level network characteristics that are robust to such sources of uncertainty.

Our own experience suggests that archaeological network analysis is a powerful tool for identifying and exploring patterns of social relations among settlements and larger areas at a variety of social and geographic scales. We have found that network methods are most powerful as a means for characterizing strong patterns using high quality and highly standardized data rather than teasing out subtle patterns or dealing with piecemeal data. As the challenges outlined above suggest, it is important to carefully consider all of the assumptions that are used to create archaeological networks as well as the numerous potential sources of error introduced by collection methods, sample size, contemporaneity, missing data, and scale.

While it may be tempting to view the challenges highlighted above as methodological shortcomings of archaeological network analyses, we would like to end with the suggestion that they may be more appropriately viewed as methodological and theoretical opportunities. For example, one might imagine that the use of archaeological materials for defining network edges is merely a surrogate for the direct measures of social interaction that more typically constitute social networks. However, the use of material-based network edges may offer an opportunity to explore the structure of material and knowledge flows in human societies that are rarely considered in more conventional social network studies (see also Knappett 2011; see in this volume Knappett 2016; Van Oyen 2016). One might also imagine that the temporal scales at which archaeologists necessarily operate—scales that are almost always super-generational—diminish the interpretive potential of

archaeological network structure. However, such scales may reveal meaningful and significant social dynamics that are undetectable at short time-scales. Accordingly, we are confident that as archaeologists continue to grapple with challenges in archaeological network analysis, we will discover new answers to old questions and new questions from old data, and ultimately, contribute to the broader interdisciplinary field of network science.

ACKNOWLEDGEMENTS

We would like to thank the editors of this volume for inviting us to participate and for providing useful comments on this chapter. The Southwest Social Networks Project is funded by the National Science Foundation, Human and Social Dynamics program (University of Arizona #0827007 and Archaeology Southwest #0827011). We would also like to thank the numerous individuals and institutions who provided the data and access to collections necessary to create the SWSN Project database.

REFERENCES

- Abbott, D. (2000). *Ceramics and Community Organization among the Hohokam*. Tucson: University of Arizona Press.
- Bernardini, W. (2007). 'Jeddito Yellow Ware and Hopi Social Networks', *Kiva*, 72:295–328.
- Bolland, J. M. (1988). 'Sorting Out Centrality: An Analysis of the Performance of Four Centrality Models in Real and Simulated Networks', *Social Networks*, 10:233–53.
- Bonacich, P. (1972). 'Factoring and Weighting Approaches to Status Scores and Clique Identification', *Journal of Mathematical Sociology*, 2:113–20.
- Bonacich, P. (1991). 'Simultaneous Group and Individual Centralities', *Social Networks*, 13:155–68.
- Borgatti, S. P. (2005). 'Centrality and Network Flow', *Social Networks*, 27:55–71.
- Borgatti, S. P., Carley, K. M., and Krackhardt, D. (2006). 'On the Robustness of Centrality Measures under Conditions of Imperfect Data', *Social Networks*, 28:124–36.
- Borgatti, S. P., and Halgin, D. S. (2011). 'On Network Theory', *Organization Science*, 22:1168–81.
- Brainerd, G. W. (1951). 'The Place of Chronological Ordering in Archaeological Analysis', *American Antiquity*, 16:301–13.
- Brewer, D. D. (2000). 'Forgetting in the Recall-based Elicitation of Personal and Social Networks', *Social Networks*, 22:29–43.
- Brughmans, T. (2010). 'Connecting the Dots: Towards Archaeological Network Analysis', *Oxford Journal of Archaeology*, 29:277–303.
- Brughmans, T. (2013). 'Thinking Through Networks: A Review of Formal Network Methods in Archaeology', *Journal of Archaeological Method and Theory*, 20:623–62.

- Carley, K. M., Dombroski, M., Tsvetovat, M., Reminga, J., and Kamneva, N. (2003). 'Destabilizing Dynamic Covert Networks', in *Proceedings of the 8th International Command and Control Research and Technology Symposium at the National Defense War College, Washington DC, Evidence Based Research, Track 7* [online publication] http://www.casos.cs.cmu.edu/publications/papers/dombroski_2003_estimatingshape.pdf.
- Carley, K. M., Lee, J., and Krackhardt, D. (2001). 'Destabilizing Networks', *Connections*, 24:31–4.
- Clark, J. J., and Lyons, P. D. (eds) (2012). *Migrants and Mounds: Classic Period Archaeology of the Lower San Pedro Valley*. Tucson: Archaeology Southwest.
- Clark, J. J., Shackley, M. S., Hill, J. B., Peebles, M. A., and Haas, W. R., Jr, (2012). *Long Distance Obsidian Circulation in the late Pre-Contact Southwest: Deviating from Distance-Decay*. Paper presented at the 77th Annual Meeting of the Society for American Archaeology, Memphis, TN.
- Costenbader, E., and Valente, T. W. (2003). 'The Stability of Centrality Measures when Networks are Sampled', *Social Networks*, 25:283–307.
- Coward, F. (2010). 'Small Worlds, Material Culture, and Ancient Near Eastern Social Networks', *Proceedings of the British Academy*, 158:453–84.
- Cowgill, G. (1990). 'Why Pearson's r is Not a Good Similarity Coefficient for Comparing Collections', *American Antiquity*, 55:512–21.
- Crown, P. L. (1994). *Ceramics and Ideology: Salado Polychrome Pottery*. Albuquerque: University of New Mexico Press.
- Deaver, W. L. (1989). 'Pottery and Other Ceramic Artifacts', in Lange, R. C. and Deaver, W. L. (eds), *The 1979–1983 Testing at Los Morteros (AZ AA:12:57 ASM), a Large Hohokam Village Site in the Tucson Basin*. Tucson: Arizona State Museum, 27–81.
- Düring, M. (2016). 'How Reliable are Centrality Measures for Data Collected from Fragmentary and Heterogeneous Historical Sources? A Case Study', in T. Brughmans, A. Collar, and F. Coward (eds), *The Connected Past: Challenges to Network Studies in Archaeology and History*. Oxford: Oxford University Press, 85–101.
- Evans, J. (1989). 'Neutron Activation Analysis and Romano-British Pottery Studies', in J. Henderson (ed.), *Scientific Analysis in Archaeology*. Oxford University Committee for Archaeology, Monograph No. 19 & UCLA Institute of Archaeology. Archaeological Research Tools 5. Oxford: Oxbow Books, 136–62.
- Freeman, L. C. (1978). 'Centrality in Social Networks Conceptual Clarification', *Social Networks*, 1:215–39.
- Galaskiewicz, J. (1991). 'Estimating Point Centrality Using Different Network Sampling Techniques', *Social Networks*, 13:347–86.
- Gjoka, M., Butts, C. T., Kurant, M., and Markopoulou, A. (2011). 'Multigraph Sampling of Online Social Networks', *IEEE Journal on Selected Areas in Communications*, 29:1893–905.
- Glowacki, D. M., and Neff, H. (eds) (2002). *Ceramic Production and Circulation in the Greater Southwest: Source Determination by INAA and Complementary Mineralogical Investigations*. Los Angeles: Cotsen Institute of Archaeology at UCLA.

- Golitko, M., Meierhoff, J., Feinman, G. M., and Williams, P. R. (2012). 'Complexities of Collapse: The Evidence of Maya Obsidian as Revealed by Social Network Graphical Analysis', *Antiquity*, 86:507–23.
- Granovetter, M. (1976). 'Network Sampling: Some First Steps', *American Journal of Sociology*, 81:1287–303.
- Hart, J. P., and Engelbrecht, W. (2012). 'Northern Iroquoian Ethnic Evolution: A Social Network Analysis', *Journal of Archaeological Method and Theory*, 19:322–49.
- Haury, E. W. (1958). 'Evidence at Point of Pines for a Prehistoric Migration from Northern Arizona', in Thompson, R. H. (ed.), *Migrations in New World Culture History*. Tucson: University of Arizona Press, 1–8.
- Heckman, R. A. (2000). 'The Tucson Basin Tradition', in R. A. Heckman, B. K. Montgomery, and S. M. Whittlesey (eds), *Prehistoric Painted Pottery of Southeastern Arizona*. Tucson: Statistical Research, 83–94.
- Hill, J. B., Clark, J. J., Doelle, W. H., and Lyons, P. D. (2004). 'Prehistoric Demography in the Southwest: Migration, Coalescence, and Hohokam Population Decline', *American Antiquity*, 69:689–716.
- Hill, J. B., Wilcox, D. R., Doelle, W. H., and Robinson, W. J. (2012). *Coalescent Communities GIS Database Version 2.0: Archaeology Southwest and Museum of Northern Arizona*. Tucson: Archaeology Southwest.
- Huntley, D. L., Fenn, T., Mills, B. J., and Habicht-Mauche, J. (2012). 'Embedded Networks? Pigments and Long-Distance Procurement Strategies in the Late Prehispanic Southwest', in L. Cordell and J. Habicht-Mauche (eds), *Potters and Communities of Practice: Glaze Paint and Polychrome Pottery in the American Southwest, A.D. 1250–1700*. Anthropological Papers of the University of Arizona. Tucson: University of Arizona Press, 8–18.
- Knappett, C. (2011). *An Archaeology of Interaction: Network Perspectives on Material Culture and Society*. Oxford: Oxford University Press.
- Knappett, C. (2016). 'Networks in Archaeology: Between Scientific Method and Humanistic Metaphor', in T. Brughmans, A. Collar, and F. Coward (eds), *The Connected Past: Challenges to Network Studies in Archaeology and History*. Oxford: Oxford University Press, 21–33.
- Kohler, T. A., and Blinman, E. (1987). 'Solving Mixture Problems in Archaeology: Analysis of Ceramic Materials for Dating and Demographic Reconstruction', *Journal of Anthropological Archaeology*, 6:1–28.
- Laumann, E. O., Marsden, P. V., and Prensky, D. (1992). 'The Boundary Specification Problem in Network Analysis', in L. C. Freeman, D. R. White, and A. K. Romney (eds), *Research Methods in Social Network Analysis*. Piscataway, NJ: Transaction Publishers, 61–87.
- Lee, S. H., Kim, P.-J., and Jeong, H. (2009). 'Statistical Properties of Sampled Networks', *Physical Review E*, 73:016102.
- Marsden, P. V. (1993). 'The Reliability of Network Density and Composition Measures', *Social Networks*, 15:399–423.
- Martyneč, R. J. (1993). 'Decorated and Intrusive Ceramics from the SCFAP Sites', in K. T. Henerdson, and R. K. Martyneč (eds), *Classic Period Occupation on the Santa*

- Cruz Flats: The Santa Cruz Flats Archaeological Project.* Flagstaff: Northland Research, Inc., 277–311.
- Mills, B. J. (2007). ‘Performing the Feast: Visual Display and Suprahousehold Commensalism in the Puebloan Southwest’, *American Antiquity*, 72:210–40.
- Mills, B. J., Clark, J. J., Peeples, M. A., Haas, W. R., Jr, Roberts, J. M., Jr, Hill, J. B., Huntley, D. L., Borck, L., Breiger, R. L., Clauset, A., and Shackley, M. S. (2013b). ‘Transformation of Social Networks in the Late Pre-Hispanic Southwest’, *Proceedings of the National Academy of Science of the United States of America*, 110:5785–90.
- Mills, B. J., Peeples, M. A., Haas, W. R., Jr, Borck, L., Clark, J. J., and Roberts, J. M., Jr, (2015). ‘Multiscalar Perspectives on Networks in the Late Prehispanic Southwest’ *American Antiquity*, 80:3–24.
- Mills, B. J., Roberts, J. M., Jr, Clark, J. J., Haas, W. R., Jr, Huntley, D. L., Peeples, M. A., Borck, L., Ryan, S. C., Trowbridge, M. A., and Breiger, R. L. (2013a). ‘The Dynamics of Social Networks in the Late Prehispanic U.S. Southwest’, in C. Knappett (ed), *New Approaches in Regional Network Analysis*. Oxford: Oxford University Press, 185–206.
- Mizoguchi, K. (2009). ‘Nodes and Edges: A Network Approach to Hierarchisation and State Formation in Japan’, *Journal of Anthropological Archaeology*, 28:14–26.
- Neuzil, A. A. (2008). *In the Aftermath of Migration: Renegotiating Ancient Identity in Southeastern Arizona*. Anthropological Papers of the University of Arizona. Tucson: University of Arizona Press.
- Newman, M. E. J. (2001). ‘Scientific Collaboration Networks. II. Shortest Paths, Weighted Networks, and Centrality’, *Physical Review E*, 64:016132.
- Opsahl, T., Agneessens, F., and Skvoretz, J. (2010). ‘Node Centrality in Weighted Networks: Generalizing Degree and Shortest Paths’, *Social Networks*, 32:245–51.
- Pailes, M. (2014). ‘Social Network Analysis of Early Classic Hohokam Group Inequality’, *American Antiquity*, 79:465–86.
- Peeples, M. A. (2011). ‘Identity and Social Transformation in the Prehispanic Cibola World: A.D. 1150–1325’, Ph.D. thesis (Arizona State University, Tempe, AZ).
- Peeples, M. A., and Haas, W. R., Jr, (2013). ‘Brokerage and Social Capital in the Prehispanic U.S. Southwest’, *American Anthropologist*, 115:232–46.
- Peeples, M. A., and Roberts, J. M., Jr, (2013). ‘To Binarize or Not to Binarize: Relational Data and the Construction of Archaeological Networks’, *Journal of Archaeological Science*, 40:3001–10.
- Rivera-Hutinel, A., Bustamante, R. O., Marín, V. H., and Medel, R. (2012). ‘Effects of Sampling Completeness on the Structure of Plant-Pollinator Networks’, *Ecology*, 93:1593–603.
- Roberts, J. M., Jr, Mills, B. J., Clark, J. J., Haas, W. R., Jr, Huntley, D. L., and Trowbridge, M. A. (2012). ‘A Method for Chronological Apportioning of Ceramic Assemblages’, *Journal of Archaeological Science*, 39:1513–20.
- Robinson, W. S. (1951). ‘A Method for Chronologically Ordering Archaeological Deposits’, *American Antiquity*, 16:293–301.
- Schachner, G., Huntley, D. L., and Duff, A. I. (2011). ‘Changes in Regional Organization and Mobility in the Zuni region of the American Southwest during the Pueblo III and IV Periods: Insights from INAA Studies’, *Journal of Archaeological Science*, 38:2261–73.

- Shackley, M. S. (2005). *Obsidian: Geology and Archaeology in the North American Southwest*. Tucson: University of Arizona Press.
- Van Oyen, A. (2016). 'Networks or Work-nets? Actor-Network Theory and Multiple Social Topologies in the Production of Roman Terra Sigillata', in T. Brughmans, A. Collar, and F. Coward (eds), *The Connected Past: Challenges to Network Studies in Archaeology and History*. Oxford: Oxford University Press, 35–56.
- Tsirogiannis, C., and Tsirogiannis, C. (2016). 'Uncovering the Hidden Routes: Algorithms for Identifying Paths and Missing Links in Trade Networks', in T. Brughmans, A. Collar, and F. Coward (eds), *The Connected Past: Challenges to Network Studies in Archaeology and History*. Oxford: Oxford University Press, 103–20.
- Wallace, H. J. (1995). 'Chronology', in H. D. Wallace (ed.), *Archaeological Investigations at Los Morteros: A Prehistoric Settlement in the Northern Tucson Basin*. Tucson: The Center for Desert Archaeology, 443–65.
- Wasserman, S., and Faust, K. (1994). *Social Network Analysis: Methods and Applications*. Cambridge: Cambridge University Press.
- Wilcox, D. R., Doelle, W. H., Hill, J. B., and Holmland, J. P. (2003). *Coalescent Communities GIS Database Version 1.0: Museum of Northern Arizona, Archaeology Southwest, Western Mapping Company*. Tucson: Archaeology Southwest.
- Woodson, M. K. (1999). 'Migrations in Late Anasazi Prehistory: The Evidence from the Goat Hill Site', *Kiva*, 65:63–84.

5

How Reliable are Centrality Measures for Data Collected from Fragmentary and Heterogeneous Historical Sources? A Case Study

Marten Düring

INTRODUCTION

In social network analysis, centrality measures are used to translate empirical and common sense observations of social behaviour into mathematical expressions. In order to assess how well an algorithm performs in conditions of imperfect data, researchers typically first select either a random or a real-world network, compute a variety of centrality measures, and declare these values to be their point of reference. In a second step, they manipulate these referential networks by adding or removing nodes or ties, again either randomly or following a set of rules. They then compare the centrality measures of the referential network to the ones gathered from the manipulated network. Borgatti et al. (2006) used this approach on a large number of randomly generated networks and found that as long as manipulations were minor (*c.*10 per cent), results remained reasonably similar to the referential networks' measures. This approach helps to shed light on the impact of false and missing data on centrality computations, and also helps us to assess the ability of these algorithms to describe social reality (as we reconstruct it) itself.

It is surprising that the effectiveness of centrality measures to accurately describe notoriously vague concepts such as ‘power’ or ‘influence’ has not been used alongside empirical observations more often (similarly: Zemljic and Hlebec 2005: 74). In this chapter, I will compare the performance of common centrality measures with the results of an in-depth reconstruction of six historical networks: in this case, support networks for persecuted Jews during the Second World War. Data were extracted from historical narratives, contemporary and

retrospective autobiographical reports, interviews, applications for remuneration, and police interrogations. These sources provide a high level of detailed contextual information about the respective ties and actions they represent.

THE CASE STUDY

It has now become common knowledge that a small minority of Jews managed to survive the Holocaust in hiding and with support from a small and diverse group of helpers.¹ Soon after the end of the Second World War, historians, sociologists, (social) psychologists, and scholars from many other disciplines began to analyse stories of help and survival and found several answers to what seemed to be the key question: ‘Why did helpers decide to help?’ A large part of the sources that are available for research were collected by the Israeli memorial and research centre Yad Vashem. The institution is most famous for awarding the title ‘Righteous among the Nations’ to individuals who were proven to have helped in a selfless manner.

Many social scientists came to the conclusion that helping behaviour was a consequence of certain common characteristics among all helpers. Samuel and Pearl Oliner argued that they were driven by an intrinsic sense of morality and altruism and that a specific form of upbringing, including strong ethical and political values, could explain their actions (Oliner and Oliner 1992). Others looked at their socio-demographic background, for example, their education and wealth (Seligmann 1992).

However, historians have shown that helpers not only differed with regard to the moral qualities but also the intensity of their actions and the ways in which they were active (Benz and Körte 2001; Benz and Wetzel 1996, 1998, 1999, 2004; Grabowski 2008; Kosmala and Schoppmann 2002; Moore 2010; Sémeral et al. 2010; Voigt 2002). Their studies confirm that helpers came from all sorts of social backgrounds, had different motives, gave a large variety of different reasons to explain their behaviour, and had varying incomes and socializations. Case studies (Beer 2010) suggest that even the self-proclaimed motives of helpers underwent processes of conscious or unconscious reinterpretation and are thus not necessarily to be trusted.

Both in Germany and in the occupied countries, helpers and refugees acted under extreme pressure in a hostile environment and had to expect to be arrested immediately after their activities attracted the attention of anyone willing to denounce them. Consequences for helpers, scope for action, and available resources, however, varied considerably between Germany and the

¹ Parts of this paragraph were taken from a previously published paper (Düring et al. 2011) with permission from the editors.

occupied zones, as well as between the various occupied zones. Probably the most important difference between Germany and the occupied zones was the absence of organizations whose infrastructures could be used to help Jews and other refugees (Moore 2010: 111f.).

My approach neglects an international comparison in favour of an in-depth analysis of network structures that emerged under similar conditions, namely in Berlin from 1942 onwards. In Germany, the vast majority of people were, or at least had to be considered as, devoted Nazis or compliant to their rules; any requests for help had therefore to be made very cautiously and based on trusted relationships. Refugees faced regular checks by police and Gestapo, first targeted at finding Jews and later at finding young men who had deserted from the Wehrmacht. In addition, they had to fear the so-called 'Greifer'; Jews who were pressured by the Gestapo to find and report other refugees and were promised freedom from persecution for themselves and their families (Tausendfreund 2006).

These dangers, together with the regime's efforts to control black markets and any other form of deviant behaviour, meant that any written account of one's activities represented a significant threat. Gestapo agents interrogated anyone they associated with support activities in order to identify all collaborators. Transcripts of these interrogations can be quite informative, although they may contain (consciously) misleading or false information. The majority of the available sources were thus produced after the war. A larger number of helpers and refugees first gave evidence of their actions in the course of applications for reparations. Detailed questionnaires asked about their political activities, experienced persecution, physical and material damage, involvement in resistance activities, religious beliefs, and an extensive résumé. They were then asked to write down their stories. Designed with refugees and resistance fighters in mind, these documents were meant to cover provable participation in anti-Nazi activities and cases of illegal expropriation by the state. They were not meant to cover the practice of help and survival. In fact, regulations at first did not even consider helping Jews to be an act of resistance against the state (Leber and Mielke 1954). Any information the applicants provided has therefore to be weighed against their interest in receiving reparation from an institution that was not necessarily acting in their best interest. In 1958, Berlin's senator for the interior, Joachim Lipschitz, brought forward an initiative to honour helpers in Berlin. Those who could demonstrate an honourable lifestyle and witnesses to their actions were granted a small pension and a public acknowledgment of their help (Riffel 2006). Following this logic, prostitutes for example, could not be honoured. Again, administrators collected reports and data about both helpers and refugees that are now available for research.

Beginning with the applications for reparation, all sources were thus produced in settings which encouraged stories of virtuous helpers, since the respective institutions explicitly ruled out acknowledgment of ambivalent or

dubious motives. Somewhat more outspoken are reports by survivors. They, of course, focus on their story of survival; their purpose is to tell their stories from their own, often limited, point of view and are therefore not without omissions, distortions, and false memories.

Oliner and Oliner showed that roughly two-thirds of all helpers whose cases were documented in the Israeli memorial site Yad Vashem responded to requests for help (Varese and Yaish 2000). The vast majority of all helpers collaborated with others in order to facilitate their support for refugees. This suggests that the decision to help was not only a question of personality but also one of social embeddedness.

In my research, I understand the decision to help and its practice as a social process. In this process, helpers typically responded to requests for help and used peers to approve of and reinforce their belief systems, which eventually led them to act differently from the majority of the society they lived in. Many of the aforementioned studies (Oliner and Oliner 1992; Seligmann 1992; Varese and Yaish 2000) aimed to measure helping behaviour both statistically and through the comparison of individual cases.

I reconstruct in a formalized, and thus comparable way, social networks between helpers and refugees in Berlin, in order to discuss their importance both for the motivation to help and the ability of refugees to sustain a life underground. Relational data is used to map the complex relations which emerged between both helpers and refugees and among helpers. All sources are heterogeneous regarding the circumstances of their creation and the purpose they once served. This also means that their view on the respective networks differs: some, like survivors' accounts, tend to be descriptions of an ego network with strong emphasis on the narrator. Others, such as historical reconstructions, aim to cover the actions of a specific group of actors. This means that the extraction of relational data is based on a fragmentary patchwork of information and not on a homogeneous dataset.

This study took into account all available and trustworthy information on explicit acts of help as described in the primary sources that were available, the network boundaries are therefore merely defined by the actors and acts of help described in the material.

All interactions between helpers and recipients of help were coded into a database which describes the practice of help and the intensity of relations between two actors. Among them are information on the specific form and endurance of help, the date of their first encounter and a rough categorization of their motives. For each class of relations and for each type of tie within this class I developed definitions accompanied by examples. This approach is widely used in qualitative data analysis and acts as a bridge between the often fuzzy primary sources and the rigid coding systems. The type of tie 'Food and food stamps' in the class 'Forms of help', for example, was defined as follows:

Food and food stamps came from many different sources. Among them were dinner invitations, donated food stamps, black market trade and forged food stamps. If the respective food/food stamps were acquired through unidentifiable black market traders, the label 'Black Market' was given. All ties which were concerned with the trade and brokerage of food/food stamps are also covered by this code. Acts of self-help such as restaurant visits or theft are not covered by this code.

This definition is accompanied by examples taken from primary sources, such as:

She soon received her food from these people and felt at home there.

(Segal 1956: 144)

But she offered to help us, we gave her our food stamps, which we had bought on the black market, she gave us a double portion for the regular price.

(Segal 1956: 122)

The giver and recipient of such an act of help have been recorded in the database and in the class 'Form of help' the value '4' which stands for 'Food/food stamps' was entered. This means that helper and recipient share a tie of the weight 4. Similarly, '2' stands for brokerage and '3' for accommodation. Each one of these ties was then further classified with regard to the time range, the duration, etc. This creates a multiplex relation between givers and recipients of help for each act of help. Actors can be connected through various forms of help at the same time. This approach makes it possible to visualize each act of help as a tie between two actors which contains information regarding the form of help but also, for example, the period in which the help was granted.

The categories of the database were developed during the analysis of four distinct support networks. Each tie has been cross-referenced with other sources and interpreted based on contextual knowledge where possible. Sometimes detailed information about, for example, the duration of an act of help could not be retrieved, in these cases the label 'unknown' was given. Acts of help which could be inferred but addressed unidentifiable actors (a fictitious example: 'and then she helped three more refugees') were acknowledged by entering ties between the helper and, for example, 'Refugee_1', 'Refugee_2', 'Refugee_3'. While without doubt controversial regarding any mathematical descriptions, this approach was appropriate for the primarily visual-exploratory analysis of the networks in this study and deemed preferable to no recognition at all.

Overall, network visualizations of relations between helpers and refugees aid the exploration of the complexity of these relations and connect the actions of individuals with the developments of larger structures. This way, the complexity of social relations changes from an obstacle to the object of research. However, the resulting relational structures can only be interpreted

in a meaningful way when considered together with the detailed information and specifics of the original sources.

METHOD

The networks I selected for this comparison differ considerably with regard to data quality and in the ways in which they were active in helping refugees. Some, like the networks surrounding Franz Kaufmann, connected hundreds of individuals and facilitated the exchange of resources through long network paths and specialized clusters within the network. Other networks were far smaller and mainly connected by a few refugees who received help from a limited number of helpers. Figure 5.2 provides a visual impression of the structural properties of these networks; Table 5.1 gives an overview of their activities.

The title of this paper asks whether centrality measures can be considered as 'reliable' indicators of influence within social networks. Any attempt to test this, of course, requires case-specific definitions of reliability, influence, and social networks. Here, I define 'influential actors' as those who were responsible for the emergence and functioning of the covert support networks I have studied and 'not influential actors' as those who played a passive part and were less involved. In order to count as influential, actors needed to share one or more of the following characteristics: they needed to have frequently initiated the helping behaviour of others, provided particularly rare resources such as forged documents, and/or frequently provided more accessible forms of help such as food or food stamps. Accordingly, the following forms of behaviour do not count as influential: merely receiving resources, or one-off or irregular acts of help with regard to accessible resources. Actors that correspond to the definition given above are to be considered influential and were marked as such in the database. With the exception of the Kaufmann network, which is far larger than any of the others, the number of these particularly influential actors is around eleven in each network. Given the nature of helping behaviour as a social practice which was highly dependent on the creation and usage of many far-reaching social ties, we may conclude that this definition of influence is close to the assumptions which underlie common centrality measures. Reliability is understood as the correlation between influence and high centrality scores calculated from the data.

Measures were then computed with the following settings: there were no restraints concerning time, form of help, etc. Relational data collected for the years 1938–1945 was aggregated. This inevitably leads to a problem of anachronism: the networks may contain paths that never existed at the same time, for example, a tie that existed in 1938 could now bridge two clusters

Table 5.1. Main characteristics of the six networks at a glance.

	Number of actors	Prevailing form of help	Times of activity	In contact with pre-existing institutions?	Detected?
Franz Kaufmann	c.400	Forged documents and food stamps	Spring 1942—Summer 1943	Protestant churches in Berlin	Yes
Luise Meier	c.150	Escapes to Switzerland	Spring 1943—Summer 1944	No	Yes
'Onkel Emil'	c.150	Accommodation, food stamps, forged documents	1938–1945	No	No
Karl Deibel	c.150	Accommodation	Spring 1943— Spring 1944	Anarchist-communist Resistance groups	Yes
Erna Segal	c.80	Accommodation, food stamps	Spring 1942–1945	No	No
Cioma Schönhaus	c.50	Accommodation, food, forged documents	Spring 1942— June 1943	No	No

Table 5.2. Actors ranked according to degree score. Influential actors are highlighted in dark grey.

Actor	Degree
Luise Meier	40
Lotte Strauss	36
Josef Höfler	34
Herbert Strauss	25
Hugo Wetzstein	17
Wilhelm Ritzi	17
Jizchak Schwersenz	11
August Sapandowski	11
Willi Vorwalder	9
Else Behrend-Rosenfeld	9

which did not exist before 1944. However, trials with networks collected in six-month steps between 1938 and 1945 showed that the resulting structures do not necessarily produce more adequate representations. This is due to the fact that many ties which were very likely to have existed were not explicitly mentioned for each and every six-month period. Centrality measures will

therefore be more likely to reveal influential actors when applied to the aggregated dataset.

Measures for the frequently used betweenness (Freeman 1977), eigenvector (Bonacich 1972), and closeness centrality (Freeman 1978) algorithms were computed as well as degree, in-degree, and out-degree. For the sake of comparison, the PageRank (Brin and Page 1998) algorithm which is not typically used for the analysis of social ties was included as well. Duplicate edges were ignored. With the exception of the computation of degree centrality, the networks were treated as directed graphs which made it possible to distinguish helpers from recipients of help. Note that this functional definition does not correspond to the distinction between people considered ‘Jews’ or ‘Aryans’ by the German legislation at the time and that ties between nodes represent acts of help and not mere acquaintance. While it would be highly desirable to produce an overlay of helping behaviour within a pre-existing social structure, the sources do not contain sufficient information to facilitate this. The degree of an actor therefore represents the number of acts of help in which he or she was involved. Directed networks help to further differentiate this: the in-degree indicates help that an actor received; their out-degree measures how many acts of help they provided.

Table 5.2 gives an example of the network surrounding Luise Meier who helped Jewish refugees to escape from Berlin to Switzerland. In this network, a total of ten actors were identified as influential. A perfect match between the concept of influence and a measurement of their degree would mean that all of these influential actors would be present in the list of the top ten highest degree scores. Table 5.2 shows the degree scores in descending order. Of the total of ten influential actors only Lotte Strauss, Luise Meier, Joseph Höfler, Jizchak Schwersenz, and Willi Vorwalder have also received the highest degree scores. This means that five influential actors have received lower degree scores than actors who were considered to be non-influential.

RESULTS

In a first step, I measured the matches between high centrality scores and my definition of influence. Table 5.3 gives an overview of the results. An example may help to read these tables: The network surrounding Luise Meier had ninety-nine actors in it. Ten of them were considered to be influential. The row ‘Degree’ indicates that for this measure, five actors were found in the top ten high scoring actors. This means that degree centrality would have identified fifty per cent of the influential actors. Overall, centrality algorithms performed very well in the case of the Schönhaus and Onkel Emil networks, with 6 and 5 scores above seventy per cent (bottom row in the Table 5.3)—I thank the reviewers for highlighting this aspect. The networks do not stand

Table 5.3. The first row contains the name of the respective network together with the total number of actors. The second row contains the number of influential actors. This is followed by the counts and percentages of the respective matches for each algorithm. In order to make comparisons easier, the column on the right indicates how often this method has yielded a match of 70% or better.

	Meier, 99	Segal, 123	Schönhaus, 76	Kaufmann, 408	Onkel Emil, 98	Deibel, 188	Match $\geq 70\%$ per algorithm (rows)
Number of influential actors	10	11	11	19	11	12	
Degree	5	5	10	7	8	5	
Match in %	50	45	90	36	72	41	2
In-Degree	5	4	8	5	7	4	
Match in %	50	36	72	26	63	33	1
Out-Degree	5	5	9	11	8	5	
Match in %	50	45	81	57	72	41	2
Betweennes	5	4	8	7	8	3	
Match in %	50	36	72	36	72	25	2
Eigenvector	6	6	5	5	10	5	
Match in %	60	54	45	26	90	41	1
PageRank	5	4	10	10	8	5	
Match in %	50	36	90	52	72	41	2
Closeness	6	6	9	6	3	5	
Match in %	60	54	81	31	27	41	1
Match $\geq 70\%$ per algorithm (columns)	0	0	6	0	5	0	

out regarding the quality of available data. They do, however, differ regarding the permanence of key actors in the two networks: both networks are characterized by one close-knit cluster in which a stable group of key actors interacted with each other for long periods of time compared to the other networks in which changing constellations of key actors collaborated across several clusters.

It becomes apparent that the matches are in most cases rather low, averaging 53 per cent across all algorithms and networks. This means that only half of the influential actors also received the highest centrality scores. Out of six case studies, no algorithm was able to top the 70 per cent benchmark more than once or twice. This threshold of 70 per cent is an arbitrary value but represents a significant match between influence and centrality which can still be considered useful in empirical research. As a first result, we may therefore conclude that for this case study, with this definition of influence and this methodological setup, centrality measures are not able to reliably identify influential actors.

But what if we stop looking for exact matches and ask how many influential actors are present in a bigger group of actors who score high in centrality? In other words, are influential actors at least likely to receive high centrality scores? I increased the size of this group of high scoring actors to twenty per cent of the total number of actors in the respective network. This number was again picked based on the pragmatic consideration that for this still rather small group of actors, it would be feasible to do additional in-depth research on individuals and their activities.

Table 5.4 shows that influential actors do indeed quite often reach the top twenty per cent highest values of most centrality measures: the match now averages 70 per cent across all algorithms and networks. This means that roughly seven out of ten influential actors can be detected by an algorithm.

Overall the match between high centrality measures and actual influence according to the definition varies considerably between 100 per cent at its best and 27 per cent at its worst. In practice, this means that no one algorithm is suitable for reliably detecting influential actors in a network.

Figure 5.1 visualizes the differences between the two tests. The matches between influential actors and highest centrality scores as computed in the first run are depicted in black. Coloured in light grey are influential actors who are part of the group with the twenty per cent highest centrality scores. Some noteworthy differences become apparent: degree centrality, the most robust and the simplest measure, yielded the best result with 81.5 per cent accuracy, followed by the out-degree measure. The in-degree measure yielded the worst match. This is not very surprising as, by definition, an outgoing tie signifies an act of help by an actor while an ingoing tie signifies help that was received: it is unlikely that influential helpers would have been the recipients of a particularly high number of acts of help. However, it remains to be seen whether this

Table 5.4. Matches between influential actors among the top 20% highest scoring actors of the respective networks.

	Meier, 99	Segal, 123	Schönhaus, 76	Kaufmann, 408	Onkel Emil, 98	Deibel, 188	Match $\geq 70\%$ per algorithm (rows)
Number of influential actors/within top 20%	10/20	11/25	11/15	19/82	11/20	12/38	
Degree	9	10	10	16	10	5	
Match in %	90	91	91	84	91	42	4
In-Degree	7	4	8	16	7	4	
Match in %	70	36	73	84	64	33	2
Out-Degree	10	10	9	16	7	5	
Match in %	100	91	82	84	64	42	3
Betweenness	7	6	8	17	7	4	
Match in %	70	55	73	89	64	33	2
Eigenvector	5	10	10	14	9	5	
Match in %	50	91	91	74	82	42	3
PageRank	8	9	10	15	7	5	
Match in %	80	82	91	78	64	42	3
Closeness	8	10	9	13	3	5	
Match in %	80	91	82	68	27	42	2
Match $\geq 70\%$ per algorithm (columns)	4	5	7	5	2	0	




Figure 5.1. Matches between influential actors and centrality scores (in per cent).

effect can be reproduced with different data. It becomes apparent that (out-) degree centrality, which only takes into account the ties of an individual actor, performed much better than the other algorithms which take into account the network structure itself: betweenness centrality refers to the entity of shortest paths within a network which go through a node, closeness centrality refers to the shortest paths between a given node and all other nodes in the network, while eigenvector and PageRank centrality both take into account the centrality of neighbouring nodes. This finding suggests that influence and acts of help to many different actors correlate.

In conclusion, we can argue that centrality algorithms can indeed yield acceptable matches within a more generous range of the top twenty per cent highest centrality scores for any given network. Chances to identify influential actors among the high scoring actors, of course, increase with the size of the respective thresholds. In practice, given the unpredictable variations in how well an algorithm performs on a single network, it would be more advisable to compute measures with several algorithms and to average their output.

Despite this encouraging overall result, we need to acknowledge that this approach will always miss a minority of influential actors with low centrality scores who were not very well connected within the network as it was extracted from the sources. This can either suggest that they fell into the

false-negative category because they did not need to be highly connected to be influential, or that knowledge about their influence was based on contextual knowledge which could not be represented in the database.

It is worthwhile to examine these actors in a little more detail and to ask why their influence has not translated into high centrality scores. In Figure 5.2, all influential actors were highlighted in their respective networks using larger node symbols. Each network had one or two actors who were of critical importance for the organization of their activities. These initiators were connected to a majority of actors in the network and therefore appear central in all spring-embedder-based visualization algorithms. With the exception of the rather small network surrounding Cioma Schönhaus, there is clearly a tendency for influential actors to be well embedded and close to the network's initiators. Still, some highly influential actors are found in peripheral positions, and numerous very well connected but not particularly influential actors are found close to the centre of the networks, in direct contact with the initiators. This raises two questions: why is it that well-connected actors in the centre who must have been involved in many acts of help turn out to be less influential than some actors in the periphery? And how could these peripheral actors be so influential?

A re-examination of the actors in question leads back to the histories of the individual networks and the people who were involved in them. Most of the influential actors who ended up in the periphery turn out to have been affiliated with other support networks which were not considered in these computations in order to not distort measures for the individual networks. From this position, they acted as brokers who facilitated the exchange of essential resources across distinct networks. The manually annotated influential actors shown in Figure 5.2 then may be considered as indicating hitherto undiscovered, yet important parts of the network in question. As such they serve as important pointers towards further analysis.

In the case of the Kaufmann network, for example, it becomes apparent that many of the less influential actors in the centre had been involved in the Bekennende Kirche, a German protestant church formed in reaction to the Nazified 'Deutsche Christen'. The network was rooted in this social milieu and many priests and church members had been involved with support activities for church members who had been classified as Jewish in the late 1930s and early 1940s. However, their willingness to help was often limited to fellow Christians, and more importantly, by the legal status of their actions. Most of them shared trusted ties with Franz Kaufmann, the initiator of the network, but chose not to get directly involved in any illegal activities. Instead they became passive bystanders or else marginally involved in the network's efforts to produce and trade forged documents and food stamps. This visualization can therefore tell us who had high potential to become involved with the network and raises the question as to why they did not.




Figure 5.2. Mapping influential actors (highlighted using larger node sizes) within their networks. Influence does not necessarily correlate with embeddedness.

DISCUSSION

The reference point of this study was the historical reconstruction of the six networks. This approach can only identify actors who made documented contributions and assess to which extent influential actors also receive high centrality scores within their respective network. Two types of actors were often missed by centrality measures: those who as friends of a friend had brokered crucial contacts with other networks and those (often refugees) who had irregular contacts with support networks and provided essential resources on only a few occasions.

The centrality algorithms which were tested here were designed with simple models of social networks in mind and were not meant to do justice to the fallibilities and inconsistencies of real-world (historical) network data. Tables 5.3 and 5.4 show that quite surprisingly, centrality algorithms performed very differently on the six networks: none of the notions behind the different algorithms seemed to be able to identify influence significantly better than another or could be linked back to the different structures of the networks. The fact that degree centrality, the most basic algorithm, performed best is another indicator of this.

We must remember that the ties that were studied here are explicitly action based: only acts of help could be extracted from the sources, mere acquaintances were not considered. Influence correlates best with a high number of acts of help to different actors. If the logic behind centrality measures cannot be directly transferred to empirical action-based networks, we can assume that there will also be problems with networks that use proxies as indicators for social relations, such as letter-writing or mutual memberships in social groups, for example.

The results of this study lead me to conclude that for data of this kind, the notion of centrality and its mathematical and visual expression is most fruitfully understood as a *potential for centrality* which was either fulfilled, or—as the above-mentioned example from the Kaufmann network shows—for one reason or another remained unfulfilled. A high centrality score in any of these networks is therefore not necessarily a reliable indicator for actual influence. Centrality measures tend to hide less-connected influential actors from us and at the same time promote non-influential actors with high centrality scores.

Still, Table 5.3 shows that we can expect that on average c.80 per cent of all influential actors can be found in the group of actors with the 20 per cent highest scores. In practice, this can be a useful way to narrow down the list of potentially influential actors and steer the allocation of research resources.

Perhaps more interesting are the contradictions which emerge when a hermeneutic analysis of primary sources is confronted with the simplistic models behind algorithmic computations. These mismatches can bring to our attention previously overlooked parts of a network or reveal an existing but unused potential to act. Such an open approach to network computation

and visualization is arguably well suited to making the most of the exploratory and hypothesis-driven work with network analysis tools and concepts.

CONCLUSION

This article tested the accuracy of centrality measures using detailed empirical knowledge of six covert support networks for persecuted Jews during the Holocaust. The in-depth historical analysis allowed us to annotate particularly influential actors and to compare this list to the output from common centrality algorithms.

Centrality measures are useful to narrow down the list of potentially influential actors in a network but will always fail to detect those c.20–30 per cent of actors whose influence does not correspond with above average connectedness. This chapter suggests that a more promising way to integrate the otherwise too simplistic centrality models in historical research: the comparison between empirical observations and computations itself and the new questions it raises is an example of how traditional historical research methods and computations can benefit from each other.

ACKNOWLEDGEMENTS

This research was made possible by generous support from the Gerda Henkel Stiftung. I would like to express my gratitude to the editors of this volume and an external reviewer for their thoughtful comments which have helped me to improve this article.

REFERENCES

- Beer, S. (2010). ‘Helene Jacobs und die anderen Deutschen: Zur Rekonstruktion von Hilfeverhalten für Juden im Nationalsozialismus’, in B. Schmidt-Lauber and G. Schwibbe, (eds), *Alterität. Göttinger kulturwissenschaftliche Studien*. Göttingen: Schmerse, 85–110.
- Benz, W., and Körte, M. (2001). *Solidarität und Hilfe für Juden während der NS-Zeit: Band 4*. Berlin: Metropol.
- Benz, W., and Wetzel, J. (1996). *Solidarität und Hilfe für Juden während der NS-Zeit: Regionalstudien 1*. Berlin: Metropol.
- Benz, W., and Wetzel, J. (1998). *Solidarität und Hilfe für Juden während der NS-Zeit: Regionalstudien 2*. Berlin: Metropol.
- Benz, W., and Wetzel, J. (1999). *Solidarität und Hilfe für Juden während der NS-Zeit: Regionalstudien 3*. Berlin: Metropol.
- Benz, W., and Wetzel, J. (2004). *Solidarität und Hilfe für Juden während der NS-Zeit: Regionalstudien 4*. Berlin: Metropol.

- Bonacich, P. (1972). 'Factoring and Weighting Approaches to Status Scores and Clique Identification', *The Journal of Mathematical Sociology*, 2:113–20.
- Borgatti, S., Carley, K., and Krackhardt, D. (2006). 'On the Robustness of Centrality Measures under Conditions of Imperfect Data', *Social Networks*, 28:124–36.
- Brin, S., and Page, L. (1998). 'The Anatomy of a Large-scale Hypertextual Web Search Engine', *Computer Networks and ISDN Systems*, 30:107–17.
- Düring, M., Bixler, M., Kronenwett, M., and Stark, M. (2011). 'VennMaker for Historians: Sources, Social Networks and Software', *REDES- Revista hispana para el análisis de redes sociales*, 21.
- Freeman, L. (1977). 'A Set of Measures of Centrality based on Betweenness', *Sociometry*, 40:35–41.
- Freeman, L. (1978). 'Centrality in Social Networks: Conceptual Clarification', *Social Networks*, 1:215–39.
- Grabowski, J. (2008). *Rescue for Money: Paid Helpers in Poland, 1939–1945*. Jerusalem: Yad Vashem.
- Leber, E., and Mielke, H. (1954). *Wegweiser zum Entschädigungsrecht*. Frankfurt am Main: Kommentator Verlag.
- Kosmala, B., and Schoppmann, C. (2002). *Solidarität und Hilfe für Juden während der NS-Zeit: Band 5*. Berlin: Metropol.
- Moore, B. (2010). *Survivors. Jewish Self-help and Rescue in Nazi-occupied Western Europe*. Oxford: Oxford University Press.
- Oliner, S., and Oliner, P. (1992). *The Altruistic Personality: Rescuers of Jews in Nazi Europe*. Touchstone.
- Riffel, D. (2006). *Unbesungene Helden: Die Ehrungsinitiative des Berliner Senats 1958 bis 1966*. Berlin: Metropol.
- Segal, E. (1956). *Unpublished Autobiographical Report in Dossier Erna Segal*. Berlin: Gedenkstätte deutscher Widerstand.
- Seligmann, A. (1992). 'An Illegal Way of Life in Nazi Germany', in J. Grenville, and R. Gross, (eds), *Leo Baeck Institute Yearbook*. Oxford: Oxford University Press, 327–62.
- Sémelin, J., Andrieu, C., and Gensburger, S. (2010). *Resisting Genocide: the Multiple Forms of Rescue*. London: C. Hurst.
- Tausendfreund, D. (2006). *Erzwungener Verrat: Jüdische Greifer im Dienst der Gestapo 1943–1945*. Techn. Univ., Diss.–Berlin, 2005. Berlin: Metropol.
- Varese, F., and Yaish, M. (2000). 'The Importance of Being Asked: The Rescue of Jews in Nazi Europe', *Rationality and Society*, 12:307–34.
- Voigt, K. (2002). *Solidarität und Hilfe für Juden während der NS-Zeit: Band 6*. Berlin: Metropol.
- Zemljič, B., and Hlebec, V. (2005). 'Reliability of Measures of Centrality and Prominence', *Social Networks*, 27:73–88.

6

Uncovering the Hidden Routes: Algorithms for Identifying Paths and Missing Links in Trade Networks

Constantinos Tsirogiannis and Christos Tsirogiannis

INTRODUCTION

A network is a simple yet powerful tool for representing a set of relations in the real world. For instance, to represent direct business relations between several people, we can sketch a network where each person is represented by a node and any two people that have done business together are connected by a link. A naive analysis of this network gives a picture of the direct connections between individuals, that is, who has done business, in person, with whom. However, for several network applications it is important to observe more complicated structures, other than the direct connections between the nodes. An example comes from applications in trade networks, where goods are exchanged between several people. In this case, it is important to keep track of the paths that specific goods have traversed in a network; in other words, we want to know the exact sequence of nodes through which a specific item was exchanged.

Unfortunately, in some studies of trade networks we may not always know the exact path that certain items followed in the network. This is frequently the case with networks that represent trade relations between sites in an earlier historical period; knowledge of the exact trade paths in such networks has not survived and only fragmentary data is available (Sindbæk 2007, 2013; see in this volume Peeples et al. 2016). The same problem also arises in modern trade networks, when the transactions involved are the result of illegal activities. Such an example is the modern trade network in illicit antiquities. During recent decades, thousands of antiquities were illegally excavated worldwide and exchanged via a global trade network. During the late 1990s and 2000s, the combined efforts of forensic archaeologists and police investigators uncovered a considerable part of the trade network that handled illicit Italian and Greek

antiquities in particular (Gill and Chippindale 2006; Gill and Tsirogiannis 2011; Godart et al. 2008). However, a large part of the activities that took place within this network remain unknown and there are some transactions that police investigations were not able to trace. Several artefacts are known by the authorities to have been traded illegally, but the specific people involved in these transactions are unknown. In most of these cases, the authorities know that, at some point in time, an artefact was in the possession of a person A, and that later the same artefact somehow came into the possession of a person B. Yet, it remains unclear what was the exact sequence of people through which this item was traded between individuals A and B. We call this sequence the **transaction path** between A and B for this artefact.

Another problem arising from the incomplete data that we may have for a trade network is that several links of the network may not be known. For example, in the context of the illicit antiquities trade network, for each person in the network it is important to know with which other individuals at least one artefact has been exchanged. Unfortunately, for some pairs of traders such direct transactions are unknown.

The aforementioned problems highlight the need for a tool that archaeologists can use to reconstruct events that could not be traced by the standard means of investigation. For the cases of networks that we mentioned above, one of the few available tools for extracting the missing paths and links is the network itself. In this chapter we present three algorithms that compute transaction paths between pairs of network nodes, and two algorithms that indicate missing links, based exclusively on the structure of the given network. To evaluate the performance of our algorithms in practice, we have constructed a simplified model of the recent trade network of Italian and Greek illicit antiquities. We use this model to test to what extent our algorithms can identify paths of transactions and links that appear in the actual network. One of the methods that we used for estimating transaction paths produced promising results. From the methods that we used for estimating missing links in the network, one method had a remarkable performance, exposing an interesting property in the network's structure.

It is important to state here that the processes that determine the course of transactions in a social network are too complex to be simulated accurately by any mathematical method. This is because there are many real-world factors that influence the decisions of the people who take part in the network. Therefore, we cannot expect that an algorithm which is based only on the structure of the trade network can correctly estimate most of the transaction paths. However, the output that is produced by such an algorithm can be used as a point of reference for further investigation; this output can be considered as an indication of the 'most probable' path that an item may have followed in the network.

The rest of this chapter is organized as follows. First, we briefly describe the network model that we use for representing the transactions in the recent

trade in Italian and Greek illicit antiquities. As mentioned, we built this model in order to evaluate experimentally our algorithms for finding paths and links in a network. We go on to describe three algorithms for estimating transaction paths in the network. We also describe the experiments that we conducted to evaluate the performance of these methods. We then present two methods for indicating missing links in the network, together with an experimental evaluation of these methods. Finally, we summarize our results and suggest directions for future research.

MODELLING A TRADE NETWORK

Preliminaries and Mathematical Definitions

To model the network of the illicit trade of antiquities, we consider two different representations. For both representations, the nodes of the network correspond to people or foundations that have participated in at least one exchange of an item, and the links in the network connect pairs of nodes who have directly exchanged at least one item, that is, without the intervention of other people. In the first representation, links are directed to indicate the direction of transactions; who sold items to whom. In this version of the network, each directed link from a node x to a node y is assigned a positive weight; this weight is equal to the number of items that are known to have been sold from node x to node y . It is possible that the same pair of nodes x and y can be connected by two links with opposite directions; this is the case if x has sold items to y , and y has also sold items to x . This representation allows us to track the transaction paths of items in the network and to highlight which links were used to carry out many transactions. We call this version of the network the **directed** network.

In the second representation of the network, links between pairs of nodes are undirected; that is, if there was any direct transaction between a pair of nodes then we connect these nodes with a unique undirected and unweighted link. Therefore, in this representation there can exist at most one link that connects the same pair of nodes. This representation is more appropriate for examining the structure of acquaintances in the network. We call this version of the network the **undirected** network.

We use V to denote the set of nodes in the network. For any node x in V , the degree of x is the number of links that are adjacent to this node. For a pair of nodes x and y for which there exists a link from x to y in the directed version, we use $w(x,y)$ to denote the weight of this link. For two nodes x and y in V , a **path** between x and y in the directed network is a chain of links such that, starting from node x , if we follow the direction of the links on this chain we

will reach node y . However, we cannot use the same path to reach node x starting from y . In fact, there can exist one or more directed paths from x to y , while there can be no such path from y to x .

Building the Network

In order to test the performance of algorithms which indicate paths or missing links in the trade network of illicit antiquities, we built a model of this network based on our knowledge from the criminological literature. We build a ‘restricted’ model, that is we construct a model of the network that does not reflect our entire knowledge of the trade in illicit antiquities; we can split the set of transactions of illicit antiquities that appear in the related literature into two subsets; a ‘base’ subset, and a ‘test’ subset. We use the transactions of only the ‘base’ subset to create the nodes and links of the restricted network model; in this way, we construct both the directed and the undirected version of this network. Then, we apply our algorithms on the restricted network to evaluate how well they can estimate the transaction paths and the links that appear in the ‘test’ subset of transactions.

In our study, we constructed a model of the network using a small but representative subset of the related literature, that is the network described in Watson and Todeschini’s book *The Medici Conspiracy* (2007). This book is a highly acclaimed resource that presents a major part of the trade network of Italian and Greek illicit antiquities during the period from 1972 to 2001. Of course, this book does not cover the entire set of recorded transactions of illicit antiquities during this period. We consider that this resource provides a nice general picture of the network, while at the same time we are aware that it is far from being complete. For the remaining part of this chapter, we refer to the network that we extracted from *The Medici Conspiracy* as the **modelled** network. In this way, we distinguish the model from the **actual** network that existed in the real world.

The network described by Watson and Todeschini has ninety-seven nodes; the directed version of the network has 181 directed links, while the undirected version has 171 links. The maximum number of links that an undirected network of n nodes can have is $n(n - 1)/2$, which is equal to 4656 links when $n = 97$. The directed version of the network can have, at most, twice as many links. This means our modelled network is relatively sparse. The network consists of two connected components, one large component that consists of ninety-two nodes and 177 directed links (167 undirected), and a very small component that consists of only five nodes and four links (both in the directed and the undirected version). The small component represents the transaction of a single artefact among nodes that seem unrelated to the main core of the network. For this reason, we have decided to focus on the main component of




Figure 6.1. The undirected version of the network based on data from Watson and Todeschini (2007). Here we see the main component of the network that consists of ninety-two nodes and 167 links.

ninety-two nodes; for the rest of this chapter we refer to this component as if it were the entire network described in *The Medici Conspiracy*. For an illustration of the undirected and directed versions of the network, see Figures 6.1 and 6.2.

The ninety-two nodes of the network correspond to individuals, but also to organizations that were involved in the illicit antiquities trade. Most of the individuals in the network can be placed in the following groups: looters, art dealers, middlemen employed by dealers, private collectors, and employees of museums. As for the nodes that represent organizations, almost all correspond either to museums and galleries, or to auction houses such as Christie's and Sotheby's. Given the distribution of links among the nodes of the network, we have observed that the nodes that correspond to art dealers have the largest number of neighbours in the network. In particular, five of the nodes that correspond to dealers are adjacent to 108 links, which is almost 65 per cent of the total number of links. Below, where we examine algorithms for identifying




Figure 6.2. The directed version of the network based on data from Watson and Todeschini (2007).

missing links in the network, we test the hypothesis that it is more likely to find missing links among pairs of nodes that have a large degree than in the rest of the network.

ESTIMATING TRANSACTION PATHS IN THE NETWORK

One of our main aims in this paper is to design algorithms that can estimate the path an item traversed between two nodes in the network. In the context of the illicit antiquity trade network, due to incomplete data, there are many cases where we know that an artefact was in the possession of one person and eventually ended up in the hands of another, but we do not know who took hold of the item in between. Returning back to the directed representation of the network, this translates to the following problem: given a start node x and

an end node y , we want to find the most ‘reasonable’ path along which an item travelled from x to y . Here we consider a path from x to y to be a sequence of directed links, starting at x and ending at y . Of course, there are many ways to define what a ‘reasonable’ path between two nodes is, if such a path exists. In our approach, we try to extract a path based on the structure of the directed network and the weights of its links; in this way we consider that the extracted path was produced according to the knowledge that we have about past transactions that took place in the network.

Properties of Transaction Paths

Before designing an algorithm that outlines transaction paths in a trade network, we need to decide what a realistic transaction path looks like. In other words, we should describe which properties should be satisfied by the structure of a transaction path.

First, we want the transaction paths to be **simple** paths: that means each path should visit a network node at most once. Hence, a simple path does not contain a cycle of links that starts and ends at the same node. On the contrary, from the police archives we do know of cases where the transaction paths of certain artefacts contained cycles (Watson and Todeschini 2007). Nevertheless, such cases are extremely rare and should not be considered as the standard situation.

Another condition is that transaction paths should consist only of nodes and directed links that already exist in the network; this means that the transaction paths should not introduce new nodes or links. Therefore, if the modelled network does not include an actual node or link that appears in a real transaction path, then this path cannot be retrieved. Finding missing links in the network (which can then be used for estimating transaction paths) is an independent problem that we examine below. According to the last condition, given a start node x and an end node y , there is a specific subset of network nodes and links that can appear in a valid transaction path from x to y . This subset of the network is the union of all possible simple paths that start from x and end at y . This subnetwork contains all the information that can be used to construct a transaction path that satisfies this last condition. We refer to this subnetwork as the **path subnetwork** from x to y . The path subnetwork from x to y can be quite complicated, and difficult to process manually, even if the distance from x to y is very small. See Figure 6.3 for an illustration of such a case. There we can observe the path subnetwork from Robert Hecht to the Getty Museum in our modelled network. The distance between this pair of nodes is equal to one (i.e. they are directly connected with a link), yet there are twenty-one other nodes that appear on simple paths between this pair.




Figure 6.3. The path subnetwork between the node that corresponds to Robert Hecht and the node that corresponds to the Getty Museum.

Description of the Algorithms

Now we describe three different algorithms that can be used for estimating transaction paths between given pairs of nodes. For all three algorithms that we consider here, the structure of the output transaction paths satisfies the two conditions presented above. The input of each of these algorithms is the directed representation of the network, and two network nodes, a start node x and an end node y . The output is a directed transaction path connecting x to y . Two of the algorithms that we examine are novel methods designed for our study, and the third algorithm is the standard method of computing the shortest path between two network nodes. We will also provide a thorough description of these methods. The description of each method consists of two parts: first, we describe formally how the method works, and second, we explain some intuitive reasons that this method seems reasonable in a social context.

The Local Expansion Method

Technical description: in this method, we expand a transaction path incrementally, by adding at each step a link of the largest possible weight, while ensuring that the path does not contain a cycle. More specifically, we begin from the starting node x , and we look for all neighbouring nodes of x . In this context, a neighbour of x is a node connected to x with a directed link whose source is x . Of all the neighbours of x , we pick the node v_1 that satisfies the following properties: (1) there exists a simple path from x to y that contains the directed link (x, v_1) , and (2) this link from x to v_1 has the maximum possible weight. After selecting v_1 in this way, we add this node and the link (x, v_1) to the transaction path. If v_1 is the same node as y , the end node of the queried path, then we stop and we output the computed transaction path. Otherwise, we continue expanding the path from v_1 in a similar manner: looking at the neighbours of v_1 , we pick the node v_2 such that there exists a simple path from x to y that includes links (x, v_1) and (v_1, v_2) , and the value $w(v_1, v_2)$ is the largest possible. We continue this process until the last node added to the path is y .

At any step of this process, to check if there is a simple path from some node v_i to y we execute a breadth-first search (Cormen et al. 2009: 594) from v_i . To make sure that the path that we expand is a valid simple path, each time that we add a node v_i to the path, we remove all the incoming links of this node. In this way, whenever we want to find out if some other node v_j has a simple path to y , we are sure that v_j does not appear in such a path. This ensures that no node appears twice in the expanded path.

Intuition: the transaction path of an item is expanded by adding at each step the most ‘probable’ connection: when a member x of the network acquires the item, this item is then forwarded to the member v with whom x has had the most transactions so far. Since we know that eventually the item came to the possession of member y , we also make sure that x always hands on the item to a member v such that it is still possible for this item to reach y from x .

This method is based on taking decisions on a local scale: each time that we expand the transaction path we only look at the connections of the last member on the path. The transaction path is mainly determined by the connections of each member independently, and not as much by the global structure of the network.

The Maximum-Weight Arborescence Method

Technical description: this method constructs a transaction path between x and y that maximizes the minimum weight that any link in the path may have. In particular, starting from x this method constructs a **maximum weight arborescence** in the network rooted at x . In graph theory terminology, an arborescence is a directed network that has a tree structure and which satisfies

the following property: among its nodes, there exists a distinct root node such that there exists a unique directed path between the root and any other node in the network (Georgiadis 2003). This implies that an arborescence does not contain cycles; all directed paths from the root to the rest of the nodes are simple paths.

The method first computes the maximum weight arborescence that has x as root, and then extracts the unique path that connects x with y in this arborescence. This path is known to have a special property: the minimum-weight link in this path has the largest possible weight compared to all other paths that lead from x to y in the network. In other words, let p be the path from x to y that we extracted from the maximum weight arborescence that has x as root. Let p' be any other directed path in the network that starts from x and ends at y . Then, the smallest weight that can be found among any link in p can never be smaller than the smallest weight that can be found among the links of p' .

Intuition: given two nodes x and y , the maximum-weight arborescence method constructs a transaction path whose connections are as ‘fat’ as possible; the minimum number of transactions that have taken place at any link in this path is the largest possible compared to any other path that connects these two nodes. Unlike the local expansion method, this method constructs paths by looking at the network on a global scale: the computed transaction path maximizes a property over all possible paths that connect the queried nodes in this network.

The Shortest Path Method

Technical Description: this method simply computes the shortest path between the given nodes x and y : a path from x to y that consists of the minimum number of links. This is a standard method used for a wide range of network applications, such as transportation and social network analysis (Fu et al. 2006).

Intuition: we use this method as a point of reference for the other two algorithms. On the one hand, we want to test a simple method that does not depend primarily on the number of past transactions that took place on the connections of the network. More than that, we want to see how complicated the actual transaction paths are in comparison to the shortest possible routes in the network.

Experimental Evaluation

We implemented the three algorithms we have described and conducted experiments to evaluate their performance. The implementation was done in the C++ programming language. The experiments were carried out using the

directed version of our modelled network and two different sets of queries. Each query is a pair of nodes: a start node and an end node of the transaction path for a real artefact that was illegally traded in the past.

The set of queries that we used contains thirteen pairs. These correspond to transaction paths of actual artefacts retrieved from several criminological sources that refer to the examined illicit trade network. We intentionally chose paths that exclusively contain nodes that appear in our modelled network. However, most of these query paths contain at least one link that does not appear in the modelled network; this is the case for eight out of thirteen paths. As mentioned earlier, if a node or link of a path is not present in the modelled network then this path cannot be retrieved from any of the algorithms that we examine. We decided to pick such paths for two reasons: first, choosing only paths that are present in the modelled network introduces a selection bias in our evaluation. Many of the paths that appear in the related literature contain links that are not described by Watson and Todeschini (2007). We want this to be also reflected in the set of paths that we consider in our evaluation. Second, we want to measure the performance of the algorithms even when they cannot predict accurately the correct transaction paths; in such a case we want to observe to what extent the computed path differs from the correct one.

To measure the difference between a computed path and the original path, we calculated a version of the **edit distance** between the two paths (Masek and Paterson 1980). In the most usual setting, the edit distance is a measure that describes how different two words are. In short, this distance is defined as the minimum number of letters that we have to add, subtract, or substitute in the first word so that it becomes identical to the other. For example, the words ‘stalker’ and ‘snake’ have an edit distance equal to three because we have to change at least three letters in the first word to turn it into the second word.

In our application, to compare two transaction paths we use a slightly different definition of the edit distance. Given the two paths, which are two ordered sequences of nodes, we define the edit distance between them to be the minimum number of nodes that we have to add, subtract, or substitute in order to make one sequence identical to the other. Hence, the transaction paths ‘Becchina → Sotheby’s → Shelby White’ and ‘Medici → Hecht → Sotheby’s → Shelby White’ have an edit distance equal to two, because we have to add one node in the first path (‘Medici’), and substitute one more (put ‘Hecht’ instead of ‘Becchina’) in order to make it identical to the second path.

To evaluate the performance of the examined algorithms, we divided the thirteen query paths into two sets. In the first set we have the five paths that appear in our modelled network. The second set contains the eight query paths for which there exists at least one link in each path that does not appear in our modelled network. For each algorithm that we examine, we feed each query pair as input to the algorithm, together with the directed version of our

Table 6.1. The performance of the three algorithms that estimate transaction paths, given the set of thirteen queried paths. ‘Set A’ refers to the set of the five transaction paths that exist in our modelled network. The performance of each algorithm appears as the percentage of the paths that the algorithm computed and which are identical to the original query paths. Set B corresponds to the set of the eight query paths that contain at least one link that is missing from the modelled network; the performance of each algorithm appears as the percentage of the paths that were computed by the algorithm and have an edit distance one from the original query path.

Method	Set A	Set B
Local expansion	3/5 (60%)	3/8 (37.5%)
Maximum-weight arborescence	2/5 (40%)	3/8 (37.5%)
Shortest path	1/5 (20%)	3/8 (37.5%)

modelled network, then we compare the path computed by the algorithm with the actual transaction path for this pair. For the first set of queried paths, we measured how many of the paths were computed exactly by each algorithm. For the second set of query paths, we measured the edit distance between the computed path and the actual path, and calculated in how many cases the computed path has an edit distance one from the actual path (which means that the returned path is almost identical to the actual one). For the results of these experiments, see Table 6.1.

Discussion

We see that for the first set of queries the local expansion method yielded better results than the two other methods: this method correctly guessed most of the paths, achieving a performance of 60 per cent. For the second set of queries, all of the algorithms had the same performance: they returned paths that are almost identical to the correct ones in three out of eight cases. Recall that for this set of queries it is impossible for the given algorithms to compute the correct path exactly. From these results, we conclude that the local expansion method has quite a good performance, when the queried path exists in the modelled network. Recall that the local expansion method expands a path by taking decisions on a local scale: the path is expanded by one node at a time, and each time the last node in the path is the one that ‘defines’ which one is the next to be added. Given that this method performs better than the other two algorithms, we consider this as an indication that, to some extent, the structure of a path in this specific trade network is influenced more by the neighbourhood of the network nodes, rather than by some global property of the network.

Moreover, the experiments showed that the three methods examined have a mediocre performance when the queried paths do not exist in the modelled network. This probably indicates the limitations of using an incomplete modelled network: if there is a lot of data missing from the network that we use for our computations, there will be many cases where it is impossible to produce a correct estimation of a transaction path. Of course, in order to provide further support to these conclusions we must conduct experiments with a larger number of paths but also use different instances of the modelled network. We could then examine how the performance of the three methods changes with respect to the amount of data appearing in the network that we use for our predictions.

Looking at the results of the above experiments, the question arises whether we could develop a method that can perform much better than the ones we tested. While this goal may not be unrealistic, we still have to consider the effect of non-determinism in the studied problem. As mentioned earlier in this chapter, the events and circumstances in the real world that influence the form of a transaction path are perhaps too complicated to be modelled by any mathematical method. In the police archives we find many cases where a transaction of an artefact took place between a totally unexpected pair of individuals. Predicting the unexpected is not what we want to do. We want to develop a method that produces the ‘expected’ path for a given case, based on the available data of past transactions in the network. The authorities can then use such a path as a point of reference, rather than a definite answer for a given case. Therefore, it would be useful to have a method that not only guesses complete paths correctly, but which also produces paths that do not differ a lot from the actual ones.

FINDING MISSING LINKS

Another goal of this work is to study algorithms that indicate links that may be missing from a given network. For the network of illicit trade of antiquities there may be pairs of traders that had direct transactions with each other, but none of these transactions were ever recorded. This translates into links that are missing from the modelled network. To tackle this problem, we use the undirected version of the modelled network to produce a set of links that do not appear in the modelled network, but which seem highly probable as existing in the actual network.

Finding links in a social network is a problem that has been widely studied. The most popular variant of the problem is, given the current state of a social network, predicting which new links will be created between members of this network in the future. This is known as the **link prediction problem**.

(Liben-Nowell and Kleinberg 2007). In our case, the context is slightly different. We do not want to predict links that will appear in the network in the future; rather, we want to estimate links that already exist in the real world network, but which have never been recorded. However, when it comes to the mathematical definition of the problem, the objective is the same: given an undirected network we want to find a set of links that do not appear in this network, but which are likely to exist according to the structure of the network.

Description of the algorithms

In the rest of this section we examine two different algorithms for finding missing links in a network. The first algorithm is based on a simple measure that has been used in several case studies on social networks—this is the so-called Jaccard index (Liben-Nowell and Kleinberg 2007). The second algorithm was developed based on our intuition of the structure of the illicit antiquities trade network. For both algorithms, the input is an undirected network and a positive integer k . The output of each algorithm is a list of k links that do not appear in the input network, but which are considered as highly probable to exist in the actual, real-world network.

The Jaccard Index Method

Technical description: in this method we measure how likely it is that a link is missing between two nodes based on their common neighbours. Consider two network nodes x and y such that: (1) they are not connected with each other, and (2) they have a common neighbour. That means there is no link that connects x with y , but there exists a node γ in the network that is connected both to x and y . Let $N(x)$ and $N(y)$ denote the set of neighbour nodes of x and y respectively. The Jaccard index of the pair is equal to:

$$J(x, y) = \frac{|N(x) \cap N(y)|}{|N(x) \cup N(y)|} \quad (\text{Eq. 6.1})$$

In other words, the Jaccard index of x and y is equal to the number of their common neighbours divided by the number of nodes that are adjacent to at least one of these two nodes.

The algorithm that we consider for computing the k missing links in the network works as follows: first, we find all pairs of nodes in the network that are not connected, but which have at least one common neighbour. For each such pair we compute the Jaccard index, and we output the k pairs that have the largest value of this index. The k returned pairs are listed in order of decreasing value of the Jaccard index. In this way, the links appear in the

output in an order that indicates how probable they are according to this method.

Intuition: the Jaccard index is a way to estimate if there is a connection between two members of the network based on the percentage of their common friends; the higher this percentage, the more likely it is that these two members know each other. This measure assumes that two members of a network are usually introduced to each other through a common acquaintance. We believe this to be a reasonable assumption for closed communities, especially where the members of the community participate in undisclosed activities; because of the nature of such activities, it seems logical that the members of the network get to meet each other via common friends.

The Central Clique Method

Technical description: in the central clique method we consider that it is more probable for two nodes to be connected if both have relatively high degrees. More specifically, let x and y be two nodes that are not directly connected with each other. Let $\min_{\deg}(x, y)$ be the minimum degree among these two nodes. That is, if x has n neighbours and y has m neighbours then $\min_{\deg}(x, y) = \min(n, m)$. We call this value the **minimum degree** of the pair. The central clique method first extracts all pairs of nodes that are not connected with each other. Then, from these pairs the algorithm returns the k pairs that have the largest minimum degrees. The returned pairs are listed in the output in decreasing order of their minimum degree values.

Intuition: the central clique method assumes that the members of the network with the largest number of connections are also connected with each other. We do not suggest that this method will produce good results for all kinds of trade networks. We only consider that this method might perform well for networks where the most important nodes tend to form a tight cluster.

Experimental Evaluation

We have implemented both described methods in C++, and we used these methods to find potential missing links on the undirected version of our modelled network. In our experiments we evaluated the absolute number of correct links that are returned by each method and the order of likelihood that each method considers for a returned link. Recall that for both methods, the links in the output are listed in decreasing order of likelihood. In our experiments, we used a slightly different version of the Jaccard index method: we disregarded links between nodes that have degree smaller than four. According to the definition of the Jaccard index, pairs of nodes with very small




Figure 6.4. The performance of the two methods that indicate missing links in the network. The horizontal axis in the graph denotes the number of links that were queried by each method in a given experiment. The vertical axis indicates the percentage of links that were found to be valid in each experiment.

degrees need to have only one common neighbour in order to score high values in this measure. We wanted to filter out these cases because we do not consider them to represent probable connections in the network.

The settings that we used for our experiments are the following. We ran both methods using the undirected version of our modelled network as input and for each method queried the k most probable links missing from this network. We conducted this experiment for all values of k in the range from 1 to 30. For each different value of k and for each method, we measured the percentage of the returned links that represent actual connections in the real-world network. For every returned link we checked if there had been at least one recorded transaction of an illicit antiquity that happened between the two nodes adjacent to this link. To check for recorded transactions we referred to a large range of sources from the related literature, including archives of dealers that were confiscated by the Italian and Greek art police squads (Watson and Todeschini 2007). For the results of this experiment, see Figure 6.4.

Discussion

Based on the results of our experiments, it is clear that the central clique method outperforms the method that is based on the Jaccard index. Out of thirty links in total, the method that computes the Jaccard index

correctly identified only five missing links. This may disprove the assumption that members of the network are more likely to collaborate if they share a large percentage of common friends. On the other hand, the central clique method shows a remarkable performance when we ask for a small number k of missing links; this performance decreases gradually as k increases.

This happens because as the value of k increases the method output also links between less important nodes of the network. For example, if we use this method to compute $k = 10$ most likely missing links, the method will return pairs of links with relatively high degrees. However, if we use the same method to compute the top $k = 20$ missing links, it will return the same ten links as before plus ten more links that connect nodes of smaller degrees. Therefore, the more we move away from the core of the known network, the less the network looks like a clique. In other words, it seems that in the core of the network almost all nodes are connected with each other.

CONCLUSIONS AND FUTURE WORK

In this chapter, we have examined different algorithms for computing transaction paths in a trade network and other algorithms for indicating missing links in the network. To evaluate the algorithms that we developed, we built a simplified model of the illegal trade network that involved the exchange of Italian and Greek antiquities. We evaluated three different algorithms that estimate transaction paths in the network based exclusively on available data about past illegal transactions. One of the algorithms produced promising results; this is the algorithm that expands transaction paths incrementally by adding at each step the link that has carried the largest number of known transactions. However, there are limitations in the performance of these algorithms. These limitations possibly have to do with the amount of data that may be missing from the model of the network. Furthermore, we tested the performance of two different algorithms that indicate potential missing links in the network. We observed that the method that tracks connections among the core nodes of the network produced high quality results.

A possible direction for future research would be to examine whether we can improve the performance of methods that construct transaction paths by incorporating extra restrictions in the network. For example, we could forbid any path to have more than a given number of nodes, since this would unreasonably increase the price of a traded item. It would also be interesting to develop methods that can produce transaction paths which contain nodes that do not appear in the known model of the network.

REFERENCES

- Cormen, T. H., Leiserson, C. E., Rivest, R. L., and Stein, C. (2009). *Introduction to Algorithms. (Third Edition)*. Cambridge, Massachusetts: MIT press.
- Fu, L., Sun, D., and Rilett, L. R. (2006). ‘Heuristic Shortest Path Algorithms for Transportation Applications: State of the Art’, *Computers & Operations Research*, 33:3324–43.
- Georgiadis, L. (2003). ‘Arborescence Optimization Problems Solvable by Edmonds’ Algorithm’, *Theoretical Computer Science*, 301:427–37.
- Gill, D. W. J., and Chippindale, C. (2006). ‘From Boston to Rome: Reflections on Returning Antiquities’, *International Journal of Cultural Property*, 13:311–31.
- Gill, D. W. J., and Tsirogiannis, C. (2011). ‘Polaroids from the Medici Dossier: Continued Sightings on the Market’, *The Journal of Art Crime*, 5:27–33.
- Godart, L., de Caro, S., and Gavrili, M. (2008). *Nostoi. Repatriated Masterpieces*. Athens: Greek Ministry of Culture.
- Liben-Nowell, D., and Kleinberg, J. (2007). ‘The Link Prediction Problem for Social Networks’, *Journal of the American Society for Information Science and Technology*, 58:1019–31.
- Masek, W. J., and Paterson, M. S. (1980). ‘A Faster Algorithm Computing String Edit Distances’, *Journal of Computer and System Sciences*, 20:18–31.
- Peeples, M. A., Mills, B. J., Haas, W. R. J., Clark, J. J., and Roberts, J. M. J. (2016). ‘Analytical Challenges for the Application of Social Network Analysis in Archaeology’, in T. Brughmans, A. Collar, and F. Coward (eds), *The Connected Past: Challenges to Network Studies in Archaeology and History*. Oxford: Oxford University Press, 59–84.
- Sindbæk, S. M. (2007). ‘The Small World of the Vikings: Networks in Early Medieval Communication and Exchange’, *Norwegian Archaeological Review*, 20:59–74.
- Sindbæk, S. M. (2013). ‘Broken links and Black Boxes: Material Affiliations and Contextual Network Synthesis in the Viking World’, in C. Knappett (ed), *Network Analysis in Archaeology: New Approaches to Regional Interaction*. Oxford: Oxford University Press, 71–94.
- Watson, P., and Todeschini, C. (2007). *The Medici Conspiracy*. New York: Public Affairs.

Part III

Challenging Network Models

Can Archaeological Models Always Fulfil our Prejudices?

Ray Rivers

INTRODUCTION

The past few years have seen a significant growth in the use of quantitative models in archaeology. Such modelling has a long, albeit uneven, history, using ideas going back to Christaller's (1933) use of Dirichlet/Voronoi tessellations of space in the early 20th century to construct 'central places'. These elementary techniques for geometrizing social relations survived as late as the 1980s, most simply supplemented or replaced by other low-technology, 'ruler and compass' methods, such as Renfrew's XTENT model (Renfrew 1975a, 1975b; Renfrew and Level 1979). By that time, archaeological modelling (e.g. Clarke 1968, 1977; Johnson 1977) was already incorporating the graph theory techniques of the social geographers, exemplified by the burgeoning network modelling of proximal point analysis (PPA) (e.g. Terrell 1977, 1986) and by the network-based spatial interaction models (SIMs) of the social planners and transport modellers (e.g. Wilson 1970; Wilson and Bennett 1986). However, even when adapted to historical processes (Rihll and Wilson 1987, 1991), such quantitative models were not part of the archaeological mainstream and limitations on computer power and cost meant that they had only restricted application.

By the 1990s, these limitations were reinforced by a post-processual critique which led to a shift in emphasis in archaeology towards viewing space as a construct of human activity, a movement from 'space to place' (Hirsch 1995). Although SIMs had led this shift by downplaying the role of geography in characterizing social interactions, the dialogue about quantitative modelling turned against its reductive nature, even to the extent of shying away from mathematical analysis (e.g. Sheppard 2001). However, cheap computer power and the ready availability of commercial software packages have led to a revival in network methods. See Lock and Pouncett (2007) for an early

overview, and for the current state of the art see Kandler and Steele (2012) and recent monographs edited by Bevan and Lake (2013) and Knappett (2013), as well as the other chapters in this book.

The main focus of this article concerns the ‘retrodictive ability’ of archaeological models. In the absence of ‘laws’ of human behaviour, there is great freedom in how to proceed. However, it is surprising that key issues, such as constraints on model morphology, sensitivity of output to input, robustness, and resilience, are often not addressed in archaeological modelling. This absence of discussion has provoked the question that provides the title of this chapter. Are we, intentionally or not, in a position to construct models that give us our desired outcomes which we can then use as a post-hoc justification of our prejudices?

This is too broad a question to answer with any generality. I shall restrict myself to discussing how we might model ‘exchange’, an important part of past processes and one most amenable to quantitative analysis, which lends itself naturally to a description in terms of ‘networks’, which provide the framework for the subsequent discussion.

The organization of this article is as follows. After a brief discussion as to what constitutes a model and the reasons why exchange networks might take the form they do, I discuss how models are tested and used. My concerns are not novel (e.g. see Beaumont 1982). As concrete examples, discussed later, I and my collaborators have worked extensively on modelling Middle Bronze Age (MBA) maritime networks (Evans et al. 2009, 2012; Knappett et al. 2008, 2011; Rivers et al. 2013a, 2013b). More recently (Rivers and Evans 2014), we have re-examined the modelling of the growth of Iron Age Greek city states as originally discussed by Rihll and Wilson (1987). This chapter permits a more general view of the problems that we had to negotiate. Further information about the models examined here is given in the companion chapter in this volume by Evans (2016).

THEORIES AND MODELS

Our models are underpinned by theoretical ideas, and in our usage ‘model’ and ‘theory’ refer to these different levels of abstraction. Theories are general concepts, within whose scaffolding a model provides a simplified version of ‘the real world’. For orientation, I follow Read (2008) in distinguishing between ‘theory models’ and ‘data models’. To be more specific, *theory models* function at an ideational level, encoding in a formal way what Leijonhufvud (1997) calls our ‘systems of beliefs about the world’. More prosaically, *data models* are validated by their ability to describe the data, but that does not mean that they have no theoretical input. I am primarily interested here in theory modelling.

A prerequisite to answering the question ‘can archaeological models always fulfil our prejudices?’ is that we know how to put our models into correspondence with the data. With the data models codifying the data and theory models encoding the underlying theory, the problem is one of comparing them. This is not straightforward, since theory model outputs are not necessarily in the form of the data model outputs, which depend on how we present or curate the data. We would not expect raw data to conform to the theory model, since the model is designed to help our understanding of how the ‘real world’ works rather than demonstrate what happens in detailed reality. Coarse-graining of the data is necessary, by which we mean intentionally disregarding specified categories or typologies that it may possess, both at the level of artefacts and at the level of relationships between sites.

Pottery data, for example, in its raw form, is micro-data, to be contrasted (Clarke 1968) with data at the meso-level (‘community’) and the macro-level (‘region’). Theory model-making is articulated through the same levels of coarse-graining, often a reflection of different levels of social organization. As a crude characterization, at the micro-level we have individual agents, at the meso-level we have agents interacting with communities, and at the macro-level we have communities interacting with communities. In what follows, I shall focus on this last level, the most appropriate for the large-scale exchange networks of greatest interest to us here.

With all these caveats, what is it that we want? Prediction or, more accurately, retrodiction, works in two ways. Most simply, beginning from known or surmised archaeological data, we try to infer further data. In particular, when the data is poor, we use these predictions to suggest how the record may be made more complete. There is, however, a separate type of prediction that concerns us in this chapter. This is to predict the theory model itself or, rather, to eliminate theory models which are incommensurate with such record as exists. This may enable us to understand better the agency behind the network formation, an important prediction in its own right, if the proposal by Rihll and Wilson is correct that:

The purpose of a good model is to formulate simple concepts and hypotheses concerning them, and to demonstrate that, despite their simplicity, they give approximate accounts of otherwise complex behaviour or phenomena. If a model ‘works’ (faithfully represents the known evidence) then it shows that the assumptions and hypotheses built into the model contribute to an explanation of the phenomena. (Rihll and Wilson 1991: 69)

However, Friedman (1953) argues that even if the model ‘works’, this may say little about the validity of the assumptions behind it:

A hypothesis is important if it ‘explains’ much by little.... To be important, therefore, a hypothesis must be descriptively false in its assumptions; it takes account of, and accounts for, none of the many other attendant circumstances,

since its very success shows them to be irrelevant for the phenomena to be explained.... Truly important and significant hypotheses will be found to have 'assumptions' that are wildly inaccurate descriptive representations of reality, and, in general, the more significant the theory, the more unrealistic the assumptions. (Friedman 1953: 14).

In our model-making, the link between understanding and prediction is potentially precarious, as Friedman suggests, since we shall see that a single model can permit very different interpretations, but our applications lead us to the former viewpoint.

AGENCY

In the natural sciences, these delineations are clear because of the role played by physical law. Unlike the sciences, society is not governed by physical laws, only the aforementioned systems of beliefs that we can sometimes crystallize into 'agency'. The antithetical position—rejected here—is to deny the existence of theories *per se*, meshing into the processual/post-processual debate. For scientists of the 20th century, this denial is particularly associated with Duhem (1974) and Quine (1953), through to Lakatos (1970) and beyond. In particular, Duhem's argument that arbitrary hypotheses can be maintained in the face of arbitrary evidence if one is prepared to adjust the system of beliefs renders prediction/retrodition impossible—one just needs a little more imagination! If we were to take this seriously, we would have to abandon our programme here, but it reminds us to be modest in the questions that we ask.

Nominally, there are two very different ways to proceed that have been adopted in the literature, most simply caricatured as looking for the most 'likely' or the most 'beneficial' networks. The former is 'passive' in that it interrogates our limited knowledge and expectations of the system and points us to the most likely network commensurate with that knowledge. The latter is 'active' in that it assumes that networks only survive because they are organized 'well', tending towards some type of 'optimal' behaviour on behalf of the individuals or communities.

Most Likely Networks

As already indicated, the 'most likely' networks are seen in an epistemic light, exploring how we make the best use of our limited knowledge of the system without invoking either agents or collective behaviour directly. They are the traditional SIMs, with strong precursors in modelling traffic and retail flows.

Networks consist of sites (nodes) connected by links (edges). A simple traffic flow network model for buses that exemplifies this approach is as

follows. Fix the sites (bus-stops) and determine the cost of travel between each of them. Fix the total number of passenger journeys and the total cost of those journeys. These comprise the totality of the inputs. We then ask what the most likely distribution of journeys between bus-stops will be that satisfies these constraints. To determine this, we make a list of all possible networks of passenger flows that are compatible with the constraints and, for each, take a picture of the network. We make a flipbook of these pictures and riffle through it to see which type of network configuration is most likely. That is our answer!

By putting each picture only *once* in the flipbook we are saying that they have equal statistical weight. This is an implementation of the ‘principle of indifference’ (also called the ‘principle of insufficient reason’), attributed to Laplace, but most likely due to Leibniz (Hacking 1975). The assumption is that not to give them equal weight would imply that I had more information about the system than I had admitted to. The information about the system is encoded in its ‘entropy’, defined by the number of questions with which we interrogate it to get a complete understanding of it. What we are doing is identifying those networks which *maximize* this entropy, subject to the constraints. Let us now suppose that each bus-stop services a local ‘population’ P . For the simple case above the outcome is a ‘gravity’ model (Wilson 1970; Ortuzar and Willumsen 1994) in which the ‘exchange’ between sites labelled 1 and 2 (the passenger flow between the sites in a day, say) is of the form:

$$F_{12} = P_1 P_2 f(r_{12}) \quad (\text{Eq. 7.1})$$

where sites 1 and 2 have populations (sizes) P_1 and P_2 respectively and are separated by distance r_{12} . F_{12} is the strength of the link between 1 and 2. $f(r_{12})$, known as the ‘deterrence’ function (but which I would rather call an ‘ease of travel’ function) describes how the ‘cost’ of travel is related to separation.

More general constraints can be applied, as we shall see later. The analogy with archaeological modelling is clear. In his seminal work on spatial analysis in archaeology, Clarke (1977: 20) suggests that this maximum entropy approach ‘represents an interesting elaboration of the missing statistical and stochastic background behind the social physics approach’, an approach that has been criticized for social gravity models as lacking a theoretical base (Jensen-Butler 1972; Olsson 1970). If we had equal costs for equal distances, $f(r)$ would fall off exponentially with r . We take it as given.

Most Beneficial Networks

Let us now consider ‘beneficial’ networks. By *beneficial* the question is, of course, beneficial to whom? This depends on the scale at which we create the model. For micro-level agent-based models (ABMs) it is the agents, which

I take to be individuals, who benefit. SIMs frequently adopt a half-way position in that individuals seek to maximize their ‘benefits’ within a given stochastic societal structure, which can be translated into community–community interaction. More conventional cost–benefit models (CBMs), like our *ariadne* model (Evans et al. 2009), most simply assume that individual behaviour can be aggregated, and collective benefits accrue as the result of this effective collective behaviour, which adjusts to improve these benefits and reduce the costs.

In fact, this dichotomy between active agents and communities ‘out there’ juggling costs and benefits and the passive observer making best use of their limited knowledge can be false. This is exemplified by the example (1) above, in which gravity models can also be understood through cost–benefit analysis in archaeological modelling (Wilson 2007) and in terms of agents in econophysics (see Piermartini and Teh (2005) for references). However, for some models the translation from one form to the other can be tortuous, if not ‘unnatural’ and, *pace* Friedman, it is often more intuitive to leave things as they stand (Wilson 1970).

THE MODEL CYCLE

The choice of model is the crucial first step in the process. Most simply, a model can conveniently be envisaged as a ‘black box’ permitting different initial settings with a fixed (small) number of knobs and dials, whose nature largely characterizes the model. To take the simile further, think of our model black box as a ‘(radio) receiver’ capable of picking up a limited frequency band, and our data model as a ‘data transmitter (radio station)’. Different ‘receivers’ then encode different agency and different implementations of that agency.

The question implicit in our analysis, as to whether we always get what we want, is now seen to be ambiguous, with the answer depending on whether we adopt a Bayesian or a frequentist approach. In the former, there is a single ‘data transmitter’, the archaeological record curated as a data model. The question then becomes one of whether there are many model receivers capable of being tuned to this station. If so, the answer is essentially ‘yes’ and we invoke the one whose agency is closest to our predilections. More typically, it may be that our receiver cannot tune into our data at all—an AM radio unable to get an FM signal—and we have to hunt to find one that does. If so, the answer is essentially ‘no’.

The frequentist question is broader, asking if we can construct a model to give agreement with chosen hypothetical data, whether in accord with the archaeological record or not? We now have many data transmitters

corresponding to hypothetical data sets, which we might think of as representing counterfactual histories of greater or lesser plausibility. The answer to our question now depends not on the number of receivers that can pick up our data station but on the number of data transmitting stations which our model receiver can pick up. In particular, we are interested in those data stations which are largely ‘unheard’, whose presence takes us towards the answer ‘no’.

Inputs

For simplicity, I am only considering networks with whole sites as their nodes, whose size is their population or resources and the exchange between them their link strength. How, then, do we proceed with our given black box? There is a problem in practice in that the inputs (initial settings and knob positions) have numerical values, as do the output dials. The question is how to impose exact mathematics on inexact systems. Unlike the case (1) above, where ‘exchange’ meant numbers of travellers, in general ‘exchange’ has many components, e.g. trade, migration, exogamy, social storage, technical innovation. This is the continual tension between the sentiment (Cameron 1963: 13), often attributed to Einstein, that ‘not everything that counts can be counted and not everything that can be counted counts’ and Lord Kelvin’s (1883) comment ‘I often say that when you can measure what you are speaking about, and express it in numbers, you know something about it; but when you cannot measure it, when you cannot express it in numbers, your knowledge is of a meagre and unsatisfactory kind.’

The compromise is to express things numerically in such a way that the actual numbers do not carry too much weight. In the absence of written records often the best that can be done is to use artefacts (largely pottery) as proxies for social exchange, coarse-grained appropriately, in terms of which we can say that the exchange between two sites varies continuously from non-existent to strong. However, theory modelling relies heavily upon aggregated categories and, for the sake of argument, we adopt single-category inputs. The numerical output values (e.g. $F_{12} = 0.67$) have no direct meaning except that a high number means stronger exchange than a lower number.

With these caveats, for a given model inputs take two forms: what I have termed, *a priori*, geophysical *settings* and sociophysical knob *calibrations*. For exchange models like ours for the Southern Aegean and Rihll and Wilson’s (RW) model for Iron Age Greece the geophysical settings primarily include the list of sites and their positions, from which can be obtained their effective separations, which reflect ease of travel between them. For some models, relative distances are enough. Beyond that the settings will include carrying capacity of sites or perhaps site populations, although often populations are part of the outputs.

As for the sociophysical knob calibrations, they include, for example, the distance scale D that can be travelled in a single journey (i.e. prevailing technology), statements about total network activity. What is surprising in practice is how few in number they are (which makes the analogy with a radio sensible). Some models, like the intervening-opportunities ‘radiation’ model of Simini et al. (2012) make a great virtue of having *no* free input parameters once the settings have been chosen. Duhem-Quine under-determinism has no traction. For data models, the situation can be different. It is only too easy to begin with a model to which, as it increasingly fails to accommodate the data, we add more and more ingredients to ‘save the phenomena’ (Duhem 1969), but that is a separate issue.

Outputs

Similar considerations apply to outputs, of which the exchange ‘flows’ F_{ij} between the sites labelled i and j (as i and j range over the network) are the most important. F_{ij} is a ‘local’ *link* quantity. Site outputs include ‘local’ *site* labels, for example, site population, site ‘busyness’, that is, the exchange flowing through the site. From these local link and site properties, we can construct outputs which reflect the importance of a site or link within the network as a whole. For sites, this includes ‘centrality’, measuring the importance of a site with respect to its neighbours, and ‘betweenness’, for both sites and links, characterizing their importance in terms of the way they connect the network as ‘bridges’ (Newman, 2010). Without attributing particular meaning to the numbers, we can form ranking tables, showing, for example, which sites are largest or busiest. There are many ways to represent these outputs but the most intuitive is to draw a network, as in the subsequent figures to this article, in which site size and link thickness reflect the desired attributes (Evans et al. 2009, 2012; Knappett et al. 2008, 2011; Rivers et al. 2013a, 2013b).

Robustness

As a further step in trying to turn the qualitative into the quantitative, we introduce the notion of ‘closeness’ in inputs and outputs for a given model. The first question we ask of the model is whether its outputs are ‘close’ if its inputs are ‘close’? If a small twitch of the knobs gives output networks that are significantly different we have a potentially serious problem. This is important because of the uncertainty in the archaeological record. Not only is the record incomplete, with sites under modern settlements and unevenness in the surviving data, but it is impossible to avoid systematic bias in its analysis. If

the model responds very differently to only slightly different inputs, there is a temptation to discount any such conclusions given their uncertainties. On the other hand if, in general, the results are insensitive to input details, our necessary ambiguity in deciding what their parameters mean becomes less important. Only models which show ‘robustness’ in this general sense, which limit the effects of one’s ignorance, are of use. Some caution is required. It is not necessarily the case that a model is robust in general, but may be only for a particular set of inputs. It may be that a strong difference in outputs for some slightly varying inputs is just what we want—the tipping point—the straw that breaks the camel’s back which leads to an anticipated system collapse, say. However, if every straw disturbs the camel, we have a silly model.

This is only part of the story. Even if there was no uncertainty in our initial settings and the choice of model was unambiguous, there is still uncertainty in the functions that define the model such as the deterrence function $f(r)$. This is the problem of ‘model discrepancy’ or ‘model inadequacy’ (Kennedy and O’Hagan 2001). Even though details of the model formulation can get washed out near transitional behaviour, such behaviour is often not present and model discrepancy becomes important. Yet again we look for robustness in the sense of lack of sensitivity to model form.

The Model Cycle

We should not generalize too much, but suppose that the data are not good enough to warrant a fully-fledged data model. Nonetheless, we interrogate the theory model with simple questions of a strongly coarse-grained nature. For example, is the model sensitive to the distance scales within the network? Is the range of predicted populations (assuming population is an output) commensurate with the known, or well-guessed, range? If the model gets such basic questions wrong, we retune the knobs and then, if that fails, change the settings insofar as we can. One possibility is that no retuning or resetting gives outputs commensurate with the data, in which case we abandon the model and look for a new one. This is a common outcome, telling us just how difficult it can be to get what we want.

However, suppose the model is good enough for us to persist. We can then use it to answer questions that we have not yet asked, for which we may or may not know the answer. At this level the best questions are, typically, what does not happen, for example, the MBA Peloponnese was never dominant in the wider Aegean and we would want our models to respect that. We then interrogate the model at finer grain. The hope is that it is commensurate with further parts of the record and makes suggestions where the record is poor. We need to know when to stop before we get poor answers, which we

inevitably will if we probe too far. In this sense, I concur with Box and Draper (1987: 424) that essentially, all models are wrong, but some are ‘useful’.

Determinism and Prediction

Finally, before we see how these ideas work for a concrete system, let us return to the question of postdictions/predictions. Our models are deterministic, insofar that they can be machine-coded. However, it is not necessarily the case that determinism means prediction.

The key question is, how good is good enough? In any ranking scheme differences between adjacent sites or links in ranking lists is often small. Can we trust small differences? Given our inability to quantify outputs reliably, the answer must be no. Thus even a deterministic output encodes ambiguity. We should thus coarse-grain the outputs to take this into account. This makes the outliers, for example, those sites which perform significantly better or worse than the rest, of particular interest, since they will survive coarse-graining best. However, this is only part of the story. Even if we know what constitutes the ‘best’ we might be happy with outcomes that are close to best, that is, outcomes that are ‘good enough’ for purpose. This ‘satisficer concept’ (Simon 1957) is how we function in everyday life and how models such as cost–benefit models actually work. They are most simply understood in terms of a ‘social potential’ (Butts 2007) or ‘network landscape’, each point of which is a network. The aim is to find the ‘lowest’ or the ‘acceptably low’ points (networks) in this multi-dimensional landscape. Monte Carlo computer programs do this by controlled guesswork, doing their best to continually move the network ‘downhill’, but not necessarily reaching the lowest point. Although *identical* inputs lead to outcomes that are different in detail, they would be expected to show the same generic behaviour. However, the differences may still be very important, exemplified by Nixon writing two speeches for the Apollo crew after the first moon landing before their return, one full of joy at their safe homecoming, the other mourning their loss. This stochastic contingency, in which the difference between two different outcomes is very small, depending on the reliability of just one or two among thousands of electronic components, should be distinguished from the unpredictably rare Black Swan events (Taleb 2010), of which the eruption of Thera is one.

MODELS FOR ARCHAEOLOGICAL NETWORKS

As I have already indicated, there is a suite of models from which to choose, and I will consider some of the most significant. See the chapter in this volume by Evans (2016) for more details.

Proximal Point Analysis and 'Intervening Opportunity' Models

Let us begin with the simplest model, proximal point analysis (PPA), our first *null* model. If archaeologists know only one network model it is likely to be PPA. It is one of a class of models that uses geographical space relatively, rather than absolutely, in specifying a site's nearest, next nearest neighbours, etc., rather than specifying how near they may be. As such, it is indifferent to the limitations of travel technology. This can make sense if the typical distances travelled in single journeys are sufficiently short with regard to the maximum feasible journeys. PPA is a particularly simple model in that it gives connectivity of any site only to a specified number k of nearest neighbours. Unsurprisingly, with its simple on-off switches for links, sensitive to the value of k , PPA is not robust unless the site data is very secure, for example, well-defined coastal sites (Terrell 1977). There is a further problem with PPA as normally presented in that the links are not directed. Once the links are made, directions are dropped. This reciprocal behaviour between sites is unrealistic. This can be remedied by retaining link direction in 'directed PPA' (DPPA) models, although the insensitivity to geography remains.

PPA and DPPA are particular examples of more general 'intervening opportunities' models (IOMs), which preserve this use of site separation ranking, rather than site separation distance. Introduced by Stouffer (1940), developed by Schneider (1959) and Wilson (1970), they were posed actively, devised to minimize travel times subject to a journey concluding at a destination point with a probability that is independent of the order in which destinations are considered, but they do permit an entropy interpretation, albeit a somewhat tortured one (Wilson 1970). In general, they provide a smoothed-off version of PPA with links varying from strong to weak. This is important, in that there has been extensive discussion of the importance of weak links in stabilizing networks and improving their functionality (Csermely 2004; Granovetter 1973, 1983). However, for the purpose of this article we restrict ourselves to PPA and DPPA models.

Generalized Gravity Models

If we were to make a continuum from models emphasizing 'place' (social networks) to models emphasizing 'space' (geographical networks), our null PPA is distinctly 'place'. The most extreme 'geographical' network would be one in which sites only interact with other sites within a fixed radius, termed a 'maximum distance network' (MDN, see Evans (this volume)). Rather than take such a simple model as a null model for 'space', we take the simple gravity

model (SGM) of (1), obtained from maximizing the network entropy subject to fixing the total network activity and the ‘cost’ of maintaining the network. This cost is encoded in the deterrence function through the effective journey distance D , such that single-journey travel over distances less than D is easily accomplished and travel over distances larger than this is difficult. Suffice to say that the smaller D , the greater the cost in sustaining the network.

One feature of PPA that is sensible to retain is the constraint it imposes on outflows. We can adopt a similar constraint here, limiting outflows according to site size, in *addition* to constraining overall activity and cost. The resulting outflow constrained gravity model (OCGM) differs strikingly from the simple gravity model in two ways. Firstly, for the SGM we would have the same exchange between sites with the same populations and separation, wherever they are in the network. This is not the case for the OCGM, for which the position in the network matters. Further, with directed links the OCGM does not show the reciprocal behaviour between sites possessed by the SGM, important when large sites interact with small.

A yet further constraint used in some models is also to limit the inflows. If we label sites such that the separation between sites i and j , populations P_i, P_j respectively, is r_{ij} , the resulting doubly constrained gravity model (DCGM) for the exchange F_{ij} between the sites takes the form (Wilson 1970):

$$F_{ij} = A_i B_j P_i P_j f(r_{ij}/D) \quad (\text{Eq. 7.2})$$

where A_i and B_i are determined by the inflow and outflow of site i . We have normalized all distances with respect to the distance scale D (the OCGM is obtained from (2) by setting the B_i to unity).

Wilson ‘retail’ model (RW model)

All the gravity models above are very rigid once the settings have been made. Other constrained entropy models which use distance rather than ranking share the same issues of rigidity to a greater and lesser extent. One success in an archaeological context is a generalized constrained entropy model, originally designed to model retail trade (see Wilson & Bennett 1986), which has been used to describe the network of city states on the Iron Age Greek mainland (Rihll and Wilson 1987, 1991) and the Bronze and Iron Age Khabur Triangle (Davis et al. 2014). The model keeps outflows constrained and introduces one further parameter, termed ‘attractiveness’. In a retail context, this allows us to interpolate between a thriving ‘High Street’ and a retail environment dominated by shopping malls. The transition to a situation in

which a few sites garner most of the ‘tribute’, if we interpret ‘exchange’ this way, is reminiscent of the XTENT model (Renfrew and Level 1979). I shall return to this later.

Cost–Benefit Models: *Ariadne*

For cost–benefit modelling, I shall limit myself to the key points of our *ariadne* model which we devised for the MBA Southern Aegean (Evans et al. 2009, 2012; Knappett et al. 2008, 2011). Such a model seems to give us more flexibility than we want, since there are several functional forms to choose in addition to the deterrence function $f(x)$ which describe the details of the costs and benefits. In reality, the situation is much better than it seems. A minimal cost–benefit exchange model would have one benefit (exchange) and one cost (sustaining the system). In the previous models, population was an input. We now treat resources (carrying capacities) as inputs and populations (i.e. the ability to exploit them) as outputs, but we would like a wide range of population sizes (outputs). The simplest way to achieve this is if the benefit arises from connecting large communities to other large communities provided travel between them is not difficult. This homophily is common to many social networks (Newman 2002; Newman and Park 2003) and is naturally implemented through a ‘gravity model’ as *input*. Since the deterrence function $f(x)$ directly defines this benefit, the sensitivity of the networks to the distance scale D is guaranteed (see Evans et al. 2009).

The ‘cost’ is related to the total population and/or the total exchange. In practice, assuming just this single benefit and a single cost seems to give networks which move from essentially no activity to star networks (see Jackson (2008) for examples of this). The resolution to this problem lies in permitting additional benefits from the exploitation of local resources, upon which the inhabitants can fall back if necessary. Under fairly general assumptions, the outcome is effectively a three-parameter model with sensitivity to the distance scale D , the relative importance of local resources to exchange, and the scale of the total costs (or benefits). The question is whether this gives us the freedom to create almost arbitrary networks.

This limited range of tuning knobs also arises for cost–benefit models, although they are more like analogue construction kits than digital black boxes. However, we can impose real limitations upon them. One trick, borrowed from condensed matter physics, applies to models describing unstable social systems near their tipping points, as happens in the transition between ‘boom’ and ‘bust’, a tendency in several models. For such precarious systems near their instabilities the model details can get washed out and the actual numbers are less important.

EXAMPLE: BRONZE AGE MARITIME NETWORKS

The appropriateness of networks for a description of maritime exchange, already mentioned in the work of Terrell (1986) and developed later by Hage and Harary for Oceania (Hage and Harary 1991, 1996), is equally applicable to the Aegean archipelago where land-based sites can behave like islands because of the difficulty of land travel. Simple network modelling has been used for the Aegean archipelagos before, notably by Broodbank (2000) to explain Cycladic exchange in the Early Bronze Age (EBA) and by Davis (1982) to understand the role of Delos in the Archaic period. The MBA networks are particularly interesting from a modelling viewpoint because they arise as social organization becomes strongly influenced by geography as marine technology improves.

The MBA Southern Aegean is largely self-contained in space and time and, in Figure 7.1, we have identified significant MBA sites (labelled 1 to 39) as the nodes for our network modelling. This is the period of high Minoan culture in which Knossos, on Northern Crete, plays a dominant role. Knossos is labelled '1' and the important gateway Thera '10'. Other sites will be identified by number when necessary, but a full list is given in Knappett et al. (2011).

The nature of the data makes it difficult to ask anything but relative questions. In fact, we can hardly claim to have a data model in the sense




Figure 7.1. Important MBA sites for the Southern Aegean, including Knossos [1] and Thera [10]. The sea journey from the North Cretan coast to Thera is approximately 100km.

above, more a sequence of statements. A very basic question is to determine which sites are the ‘central’ sites in the network, for which we have some archaeological understanding. A second question concerns the betweenness of the Northern Crete city states and the Southern Cyclades, particularly the island of Thera (10) that looks to be a key gateway between the two and hence onward to the mainland. Given the importance of Knossos and the rich archaeological record at Thera, we would be inclined to disregard any model which did not give importance to the Northern Crete palatial sites and highlight their connection to Thera.

There is a more dramatic test we can make. In that period, Thera underwent a powerful volcanic eruption that rendered it uninhabitable, as well as imposing large-scale damage on the region. Examining Figure 7.1 suggests that eliminating Thera would cause serious damage to the network. Yet, by all accounts, exchange recovered rapidly across the region. A good test of our models is to see what happens to the network on removing Thera and whether it shows the observed resilience. We have discussed the eruption in Knappett et al. (2011) and references can be found there. These questions, together with others along the line of what does not happen, may seem a meagre haul, but we have already intimated that models with ‘most likely’ outcomes have very few adjustable parameters and so, already, they are strongly predictive.

Settings

All our models have some settings in common. The most important of these concerns typical intersite ‘separation’. This involves both land and sea travel, which has to negotiate headlands. We find (Rivers et al. 2013a, 2013b) that, for journeys of 100 km or less, there are four regional clusters, the Cyclades, Crete, the Peloponnese, and the Dodecanese. As we increase journey distance these clusters increasingly connect and, at separations of 120 km or so we move from a fragmented network to one that can be traversed easily, albeit in sequential steps.

The second setting concerns the size of the sites. In some models (not the RW model), site population is an input. In others, population is an output and the input is the carrying capacity of the site, that is, local resources. For simplicity, whichever is the case, we take sites to be small, medium, or large. As long as there is a reasonable ratio between them, outcomes are not significantly different.

To show how model-making works I restrict myself to the PPA, DPPA, and SG, OCG, DCG, RW, and *ariadne* models. Their successes and failures follow. These are enough to enable preliminary conclusions about the utility of network modelling in the MBA.




Figure 7.2. Testing the sensitivity of the simple gravity model (left) and the doubly constrained gravity model (right) to the distance scale D for single journeys. For clarity, links are included (with equal weight) only if they exceed a minimal threshold. Distances of 100km and more are necessary to connect Crete, the Cyclades, and the Peloponnese and Dodecanese mainlands. The networks in the figure should be compared to the map of Fig. 7.1. The Cyclades, Crete, and the mainlands are immediately identifiable.

Sensitivity to distance scales: there is no record of an active regional exchange network in the EBA because the marine technology of the time, that is the oar, only permitted distances significantly smaller than the 100 km distance scale of the network. Even with 60 km trips, the network would have remained largely unconnected by short paths at least. However, with a marine technology improving over time, by the MBA the appearance of the sail means that distances of 100 km and more are possible and, unsurprisingly, a thriving maritime network can, and does, develop. We demand this of our models. In Figure 7.2 we show the networks for the SGM and the DCGM models for the case $D = 60\text{ km}$ and $D = 100\text{ km}$.

It is very clear that the SGM shows how difficult it is to establish a network if only short distance travel is possible, whereas the DCGM provides a network even in these circumstances. In fact, this is largely true of any models with constrained outputs, including PPA and the RW model. This is because peripheral sites are required to establish exchange with sites at whatever distance, even though they are difficult to reach, just to have the correct output. Only the SGM and *ariadne* pass this distance test (although the singly constrained gravity model is not quite as bad as the doubly constrained). Our

ariadne model has been discussed in such detail elsewhere that we shall not demonstrate its sensitivity to distance scales here.

I have already argued for directed links. However, the SGM fails this test (as does PPA) in having undirected, and thereby reciprocal, links. I have also argued for links of variable strength and especially for the presence of weak links. This eliminates PPA and DPPA.

The criteria so far are largely general. Finally, to bring the models into closer correspondence with the record, we ask that they give due importance to the Northern Cretan palatial sites (Knossos, if possible) and to the links to the Cyclades through Thera as realized through the link and site weights and link and site betweenness. Some caution is needed since we are using the simplest definitions. Because of the on-off nature of PPA links, the question does not make good sense for them. There is also the issue of the level of ‘attractiveness’ in the RW model. Our conclusions are that the RW model does not easily fit our criteria for Cretan–Theran linkages run across the sensible parameter range. Whatever else, for an extensive range of parameter choices its outcomes have little to do with Northern Crete and Thera, with Phylakopi being the dominant node. For *ariadne*, we are not obliged to have strong links between Northern Crete and Thera with high betweenness, but it arises naturally much of the time (remember that our outcomes are stochastic). However, a key ingredient is that the network landscape is very ‘mountainous’; we have to tread a delicate balance between ‘boom’ and ‘bust’, so there is much less freedom than we might have thought.

In summary, with the partial exception of the SGM and *ariadne*, the other models have general flaws or do not even attempt to match the simplest data. Henceforth we restrict ourselves to the SGM and *ariadne*. If we are going to be able to get what we want, it will be here.

Eruption of Thera: to show the difference between the SGM and *ariadne* further, we need some supplementary data to test the models against, provided by the eruption of Thera and its aftermath. Despite the violence of the destruction and the removal of Thera as the gateway from Northern Crete to the Cyclades and beyond, it is argued that ‘civilization’ carries on much as before, to the extent that Late Minoan IB has often been seen as the acme of Minoan culture (e.g. see Cadogan 1976).

When it comes to comparing the SGM and *ariadne*, there is no contest! For the SGM the removal of Thera is achieved by the simple expedient of erasing the links between Thera and the rest of the network. This is in sharp distinction to the networks of the *ariadne* model shown in Figure 7.3, in which the removal of Thera leads to strong rearrangement of the network exchange.

In particular, it is easy to find strong links to the northeastern part of the network. Some caution is necessary in how we interpret the individual networks, because of the stochastic nature of the outputs (for which these are only one set). However, in a dramatic way, they motivate observations like those of




Figure 7.3. Exemplary *ariadne* networks for $D = 100$ km before (top) and after (bottom) the eruption of Thera. The networks have sites labelled (by size) by rank and links (by thickness) by exchange flow. The shift in the pattern of exchange after the eruption is very striking, with an emphasis on the northeast of the network, for which there is archaeological evidence for post-eruption activity.

Davis (1980: 336) and Cummer and Schofield (1984), who suggest that ‘the evidence points, if anything, to an increase in Minoan trading activity in LM IB, particularly in our excavations at Ayia Irini, Keos [site 14] where we literally had thousands of LM IB vases imported from outside’ (Davis 1980: 336).

This is, in fact, just what we want, in that *ariadne*, but only *ariadne* from the wide range of models that we have considered, enables us to make sense of the archaeological record. From a Bayesian viewpoint, the patterns of the outputs of the different model types are so distinctive that we only have one known model receiver that can tune in the data transmitter. From a frequentist position, these same patterns show that there are many counterfactual data transmitters that are unheard. From both viewpoints, the answer to the question is a qualified no!

EXAMPLE: IRON AGE GREEK CITY STATES

In contrast, I shall briefly reconsider the very different environment of the formation of city states in Iron Age Greece, which provided the original application of the Wilson retail model to archaeology (Rihll and Wilson 1987, 1991). The aim was to understand the rise of Thebes, Corinth, Athens, and other major sites, as neighbouring communities relinquished power to them, a combination of urbanization and synoikism (see Fig. 7.4 for the mainland sites considered by Rihll and Wilson in their analysis).

The typical distance between a site and its nearest neighbour is 5 km. In contrast to the MBA Aegean, geography is now less important in an environment in which distances between communities are small compared to the available ‘technology’ (walking, riding, portering) for which $\underline{D} \geq 15$ km. We follow Rihll and Wilson (1987) in giving the sites equal carrying capacity. Nonetheless, it is clear, yet again, that our null PPA and SGM cannot describe the aggregation of influence that we require. With cost–benefit models like *ariadne*, we have a different problem. The network landscape is now very flat. With no critical distance scale, there are only limited penalties in choosing between a wide range of ‘good enough’ site links. As a result, we have large fluctuations, making *ariadne* totally ineffectual at describing urbanization in this city state context (Rivers and Evans 2014).

Although, in this case, we do not have a dramatic event like the eruption of Thera to test the models against, the main result is that, from our suite of models, good agreement with the archaeological record (in the sense of which sites emerge as dominant) is only achievable with the retail model. However, this leads to the question of model discrepancy that we did not have to address for the MBA Aegean; the nature of the deterrence function $f(r)$. Not only did




Figure 7.4. The 109 archaic settlements based on figure 1 of Rihll and Wilson (1987) and using the same labels. Site 25 is Thebes, site 70 is Athens and site 82 is Corinth. See Evans (2013) for details and explicit data. The typical separation between a site and its nearest neighbour is approximately 5km.

the MBA network landscape constrain options but the matching of technology and geography at $D \approx 100$ km in the physical landscape meant that any reasonable deterrence function with this distance scale would give essentially the same results.

This is not the case here but not to the extent of disabling the model. In Rivers and Evans (2014), we contrasted the results of Rihll and Wilson, who adopted an $f(r)$ that could be interpreted as assuming equal ‘effort’ for equal distance, to those of an $f(r)$ that did not penalize relatively short distances but strongly cut off large distances, as one might attribute to a walking/portering culture on small distance scales. The results for such a choice were less in accord with the detailed archaeological record than those of Rihll and Wilson, whose postdictions had been remarkably successful.

This takes us back to our earlier comments on the need to coarse-grain both the data and the theory models. Once $f(r)$ and the distance scale D have been chosen, the Wilson retail model has only two free parameters.

A two-parameter fit to more than 100 sites will not, by definition, be a very good match. In many cases, site data is sufficiently poor that this is not an issue. However, we do have a list of 10–20 sites, say, whose significance we understand well in the context of this urbanization, still many more than we would expect the model to be able to characterize well. The resolution is to coarse-grain the data, most simply by increasing the ‘pixel’ size of the map. That is, although with the alternative choice of $f(r)$, it was not possible to single out key sites with the same success rate (e.g. Athens), sites in their vicinity would be highlighted.

Thebes provides an important exception to this. The significance of Thebes in this period is not in dispute and it duly appears as a major site in the Rihll and Wilson analysis. However, with a change in deterrence function, it disappears from the output, but with no neighbouring site taking up an equivalent position (Rivers and Evans 2014). This raises an interesting question. Statistically, we can get what we want on the broad view at an appropriate level of coarse-graining but, historically, Thebes has a weight beyond its statistical significance. We could argue that the requirement that Thebes be important forces us to choose the deterrence function of Rihll and Wilson, even though it seems less plausible, in agreement with Friedman (1953), as cited earlier. On the other hand, we can argue that the Rihll and Wilson result is, accidentally, too good to be true. The coarse-grained ability to comply broadly with the archaeological record is seen as a statement that the ideas behind the retail model have a broader remit, but we never expected them to be sufficient.

With all these caveats, the situation is not that different from the MBA Aegean. Yet again the patterns of the outputs of the different model types are so distinctive that we only have one known model receiver that can tune in the data transmitter, which patterns show that there are many counterfactual data transmitters that are unheard. Again the answer to the question is a qualified no!

SUMMARY

I have argued that the question posed in the title of this article is more subtle than it first appears. Nonetheless, for the two examples that I have considered, the exchange networks of the MBA Aegean and Iron Age mainland Greece, the answer is a qualified ‘no’, however the question is posed. Insofar as what we want is agreement with the archaeological record, we have been able to find models to achieve that in a ‘natural’ way in both systems. This was not inevitable. The former was posed in terms of costs and benefits of a network that sprang into existence when maritime technology was

sufficient to bridge the distances to enable the network to thrive, the latter in terms of making the best use of constrained knowledge or, equivalently, the benefits of aggregation of power. In that, we concur with the sentiment cited earlier in Rihll and Wilson (1991) that the ‘working’ of a model shows that the assumptions and hypotheses built into it contribute to an explanation of the data.

The models above, and those which we have only mentioned in passing, such as intervening opportunity models, exhaust our current sense of agency. New models will appear and new hypothetical data sets will be possible outcomes but one cannot, out of the blue, put together arbitrary numbers of constraints to create multi-parameter models without further qualifications of agency to illuminate them. The hope that we can necessarily get what we want is like taking a book out of Borges’ ‘Library of Babel’ and hunting for a meaningful statement.

REFERENCES

- Beaumont, J. R. (1982). ‘Mathematical Programming in Human Geography’, *Mathematical Social Science*, 2:213–43.
- Bevan, A., and Lake, M. (eds) (2013). *Computational Approaches to Archaeological Spaces*. Institute of Archaeology Publications, University College London. Walnut Creek, CA: Left Coast Press.
- Box, G. E. P., and Draper, N. R. (1987). *Empirical Model-Building and Response Surfaces*. New York: Wiley.
- Broodbank, C. (2000). *An Island Archaeology of the Early Cyclades*. Cambridge: Cambridge University Press.
- Butts, C. (2007). ‘Statistical Mechanical Models for Social Systems’, in A. Bejan and G. Merxx (eds), *Constructal Theory of Social Dynamics*. New York: Springer, 197–223.
- Cadogan, G. (1976). *The Palaces of Minoan Crete*. London: Routledge.
- Cameron, W. B. (1963). *Informal Sociology, a Casual Introduction to Sociological Thinking*. New York: Random House.
- Christaller, W. (1933). *Die zentralen Orte in Suddeutschland*. Jena: Gustav Fischer.
- Clarke, D. L. (1968). *Analytical Archaeology*. London: Methuen.
- Clarke, D. L. (1977). ‘Spatial Information in Archaeology’, in D. L. Clarke (ed), *Spatial Archaeology*. London: Academic Press, 1–32.
- Csermely, P. (2004). ‘Strong Links are Important, but Weak Links Stabilize Them’, *Trends in Biochemical Science*, 29:331–4.
- Cummer, W. W., and Schofield, E. (1984). *Keos III. Ayia Irini: House A*. Mainz: Philipp von Zabern.
- Davis, J. L. (1980). ‘Comment in Discussion of Session 7A’, in C. Doumas (ed), *Thera and the Aegean World. Volume 2: Proceedings of the Second International Scientific Congress, Santorini, Greece, August 1978*. London: Thera Foundation.

- Davis, J. L. (1982). 'Thoughts on Prehistoric and Archaic Delos', *Temple University Aegean Symposium*, 7:23–33.
- Davis, T., Fry, H., Wilson, A., Palmisano, A., Altawee, M., and Radner, K. (2014). 'Application of an Energy Maximising and Dynamics Model for Understanding Settlement Structure: The Khabur Triangle in the Middle Bronze and Iron Ages', *Journal of Archaeological Science*, 43:143–54.
- Duhem, P. (1969). *To Save the Phenomena, an Essay on the Idea of Physical Theory from Plato to Galileo*. Chicago: University of Chicago Press.
- Duhem, P. (1974). *The Aim and Structure of Physical Theory*. New York: Atheneum. [originally published 1906]
- Evans, T. (2013). 'Locations of 108 Archaic Greek settlements used in Rihll and Wilson 1987', *figshare.com* [online repository] <http://dx.doi.org/10.6084/m9.figshare.97395> (accessed 24 October 2014).
- Evans, T. (2016). 'Which Network Model Should I Use? Towards a Quantitative Comparison of Spatial Network Models in Archaeology', in T. Brughmans, A. Collar, and F. Coward (eds), *The Connected Past: Challenges to Network Studies in Archaeology and History*. Oxford: Oxford University Press, 149–73.
- Evans, T., Knappett, C., and Rivers, R. (2009). 'Using Statistical Physics to Understand Relational Space: A Case Study from Mediterranean Prehistory', in D. Lane, D. Pumain, S. van der Leeuw, and G. West (eds), *Complexity Perspectives on Innovation and Social Change*. Berlin: Springer Methodos Series 7:451–79.
- Evans, T., Rivers, R., and Knappett, C. (2012). 'Interactions in Space for Archaeological Models', *Advances in Complex Systems*, 15:1150009.
- Friedman, M., (1953). *Essays in Positive Economics*. Chicago: University of Chicago Press.
- Granovetter, M. S. (1973). 'The Strength of Weak Ties', *American Journal of Sociology*, 78:1360–80.
- Granovetter, M. S. (1983). 'The Strength of Weak Ties: A Network Theory Revisited', *Sociological Theory*, 1:201–33.
- Hacking, I. (1975). *The Emergence of Probability*. Cambridge: Cambridge University Press.
- Hage, P., and Harary, F. (1991). *Exchange in Oceania: a Graph Theoretic Analysis*. Oxford: Clarendon Press.
- Hage, P., and Harary, F. (1996). *Island Networks: Communication, Kinship and Classification Structures in Oceania*. Cambridge: Cambridge University Press.
- Hirsch, E. (1995). 'Landscape: between Place and Space', in E. Hirsch and M. O'Hanlon (eds), *The Anthropology of Landscape: Perspectives on Place and Space*. Oxford: Clarendon Press, 1–30.
- Jackson, M. O. (2008). *Social and Economic Networks*. Princeton, NJ: Princeton University Press.
- Jensen-Butler, C. (1972). 'Gravity Models as Planning Tools: A Review of Theoretical and Operational Problems', *Geografiska Annaler. Series B, Human Geography*, 54:68–78.
- Johnson, G. A. (1977). 'Aspects of Regional Analysis in Archaeology', *Annual Review of Anthropology*, 6:479–508.

- Kandler, A., and Steele, J. (eds) (2012). *Topical Issue on Cultural Evolution, Advances in Complex Systems*, 15(1–2).
- Lord Kelvin (William Thomson) (1883). *Lecture to the Institution of Civil Engineers*, 3 May 1883.
- Kennedy, M. C., and O'Hagan, A. (2001). 'Bayesian Calibration of Computer Models', *Journal of the Royal Statistical Society B*, 63:425–64.
- Knappett, C. (ed.) (2013). *Network Analysis in Archaeology*. Oxford: Oxford University Press.
- Knappett, C., Evans, T., and Rivers, R. (2008). 'Modelling Maritime Interaction in the Aegean Bronze Age', *Antiquity*, 82:1009–24.
- Knappett, C., Evans, T., and Rivers, R. (2011). 'The Thera Eruption and Minoan Palatial Collapse: New Interpretations Gained from Modelling the Maritime Network', *Antiquity*, 85:1008–23.
- Lakatos, I. (1970). 'Falsification and the Methodology of Scientific Research Programmes', in I. Lakatos and A. Musgrave (eds), *Criticism and the Growth of knowledge*. Cambridge: Cambridge University Press, 91–196.
- Leijonhufvud, A. (1997). 'Models and Theories', *Journal of Economic Methodology*, 4:193–8.
- Lock, G., and Pounsett, J. (2007). 'Network Analysis in Archaeology. An Introduction to Network Analysis', in J. T. Clark and E. Hagemeister (eds), *Digital Discovery. Exploring New Frontiers in Human Heritage. CAA 2006 Computer Applications and Quantitative Methods in Archaeology*. Budapest, 123–54.
- Newman, M. (2002). 'Assortative Mixing in Networks', *Physical Review Letters*, 89:208701.
- Newman, M. (2010). *Networks: An Introduction*. Oxford: Oxford University Press.
- Newman, M. and Park, J. (2003). 'Why Social Networks are Different from Other Types of Networks', *Physical Review E*, 68:036122.
- Olsson, G. (1970). 'Explanation, Prediction and Meaning Variance; an Assessment of Distance Interaction Models', *Economic Geography*, 46(supplement):223–33.
- Ortuzar, J. de D., and Willumsen, L. G. (1994). *Modelling Transport*. New York: Wiley.
- Piermartini, R. A., and Teh, R. (2005). 'Demystifying Modelling Methods for Trade Policy', *WTO Discussion Papers* 10. Lausanne: WTO Publications.
- Quine, W. V. O. (1953). 'Two Dogmas of Empiricism', in W. V. O. Quine (ed), *From a Logical Point of View. Nine Logico-Philosophical Essays*. Cambridge, Massachusetts: Harvard University Press: 20–46.
- Read, D. W. (2008). 'A Formal Explanation of Formal Explanation', *Structure and Dynamics*, 3:1–16.
- Renfrew C. (1975a). 'Trade as Action at a Distance: Questions of Interaction and Communication', in J. A. Sabloff and C. C. Lamberg-Karlovsky (eds), *Ancient Civilization and Trade*. Albuquerque: University of New Mexico Press, 3–59.
- Renfrew C. (1975b). 'Alternative Models for Exchange and Spatial Distribution', in T. K. Earle and J. E. Ericson (eds), *Exchange Systems in Prehistory*. New York: Academic, 71–90.
- Renfrew, C., and Level, E. V. (1979). 'Exploring Dominance: Predicting Polities from Centres', in C. Renfrew and K. L. Cooke (eds), *Transformations: Mathematical Approaches to Culture Change*. London: Academic Press, 145–67.

- Rihll, T. E., and Wilson, A. G. (1987). 'Spatial Interaction and Structural Models in Historical Analysis: Some Possibilities and an Example', *Histoire & Mesure*, 2:5–32.
- Rihll, T. E., and Wilson, A. G. (1991). 'Modelling Settlement Structures in Ancient Greece: New Approaches to the Polis: I', in J. Rich and A. Wallace-Hadrill (eds), *City and Country in the Ancient World*. London: Routledge, 59–95.
- Rivers, R. and Evans, T. (2014). 'New Approaches to Archaic Greek Settlement Structure', *Les Nouvelles de l'archéologie*, 135:21–7.
- Rivers, R., Knappett, C., and Evans, T. (2013a). 'Network Models and Archaeological Spaces', in A. Bevan and M. Lake (eds), *Computational Approaches to Archaeological Spaces*. Walnut Creek, CA: Left Coast Press, 99–126.
- Rivers, R., Knappett, C., and Evans, T. (2013b). 'What Makes a Site Important? Centrality, Gateways and Gravity', in C. Knappett (ed), *Network Analysis in Archaeology*. Oxford: Oxford University Press, 125–50.
- Schneider, M. (1959). 'Gravity Models and Trip Distribution Theory', *Papers of the Regional Science Association*, 5:51–8.
- Sheppard, E. (2001). 'Quantitative Geography; Representations, Practices and Possibilities', *Environment and Planning D: Society and Space*, 19:535–54.
- Simini, F., González, M. C., Maritan, A. and Barabási, A.-L. (2012). 'A Universal Model for Mobility and Migration Patterns', *Nature*, 484: 96–100.
- Simon, H. (1957). *Models of Man—Social and Rational*. New York: John Wiley and Sons.
- Stouffer, S. A. (1940). 'Intervening Opportunities: A Theory Relating Mobility and Distance', *American Sociology Review*, 5:845–67.
- Taleb, N. N. (2010). *The Black Swan* (2nd edn). London: Penguin.
- Terrell, J. (1977). 'Language and Material Culture on the Sepik Coast of Papua New Guinea: Using Social Network Analysis to Simulate, Graph, Identify, and Analyze Social and Cultural Boundaries between Communities', *The Journal of Island and Coastal Archaeology*, 5:3–32.
- Terrell, J. (1986). 'Prehistory in the Pacific Islands'. Cambridge: Cambridge University Press.
- Wilson, A. G. (1970). *Entropy in Urban and Regional Modelling*. London: Pion.
- Wilson, A. G., and Bennett, R. J. (1986). *Mathematical Methods in Human Geography and Planning*. New York: Wiley-Blackwell.
- Wilson, L. (2007). 'Understanding Prehistoric Lithic Raw Material Selection: Application of a Gravity Model', *Journal of Archaeological Method and Theory*, 14:388–411.

8

Which Network Model Should I Use? Towards a Quantitative Comparison of Spatial Network Models in Archaeology

Tim Evans

INTRODUCTION

Archaeology can be ‘site centric’. Much of the primary evidence comes from excavations based on a single site so naturally the primary sources for archaeological information are organized by site. This is a great help when establishing intra-site links, be they local spatial relationships which may help reveal functions of buildings on a site, or temporal ones, perhaps how different institutions waxed and waned within a society.

However, this organization of the primary evidence inhibits comparisons between sites. The regional and global interactions of each site must be deduced by secondary work, comparing information from a range of primary sources with their differing protocols. Yet deducing these wider relationships from finds is one of the key goals of archaeological research as only by understanding societies at all scales can we get a proper view of how society functions. In this sense, archaeologists, and social science in general, have long appreciated that societies are complex systems, with some coherent large-scale phenomena emerging from microscopic interaction, a language that physical scientists have only articulated over the last couple of decades; for example, see Ball (2004) or Lane et al. (2009).

Archaeology meets this challenge with several well-developed approaches. Some are rooted in physical science, such as through the chemical analysis of materials. Others are the product of human expertise, as when styles of product are compared across sites. There are efforts to produce secondary regional catalogues through human analysis of the primary sources, for instance see Mills et al. (2013), Sindbæk (2007), Terrell (2010) and chapter 4 of this volume (Peeples et al. 2016) for recent examples. Yet, there remain limitations.

Chemical analysis may reveal the sources of materials but not the paths used for their transfer. Stylistic analysis may be subject to unquantifiable bias. A systematic database from primary sources may be too costly to construct. Even with such a database, there is then too much information and we have to pull out the key patterns, to simplify the information into the important parts in order for us to understand what the data is telling us. Finally, the evidence available may just be too limited to make the deductions we hope for.

To address the challenge of interactions in archaeology it is useful to call upon the natural language of relationships, that of complex networks. This is because the edges (ties, links, ...) of a network represent bilateral interactions between vertices (nodes, actors, ...). While networks have enjoyed a surge of recent interest (Evans 2004), they have a long history in both social science and mathematics which suggests they have a significant and enduring role to play even after we discount any fads and fashions. I will look at models which produce networks reflecting the spatial arrangement of sites as these are one of the most important ways networks can be used in archaeology.

Unfortunately, there are many such spatial network models (e.g. see Evans et al. 2012; Rivers et al. 2013), often with several parameters. So the question I address here is how to choose a model from the sea of possibilities. I will not provide a definitive answer to this question as I believe this depends on the context. Rather, I will define an approach which defines families of similar models, all producing similar networks, so reducing the number of modelling possibilities facing a researcher. Hopefully, each family of network models will be linked to different types of problem. Further, by looking at networks produced from models in the same group, robustness of any answers obtained may be tested. That is, if models from the same family show the same property, we can be confident that we have identified a key feature of our problem. In this way, my approach is one response to the questions raised by Ray Rivers (2016) in the present volume. The methods suggested here are therefore part of a wider discussion of robustness in archaeological methods, an issue raised in different contexts in this volume, for instance in the chapters by Peeples et al. (2016) and Düring (2016).

SPATIAL MODELS FOR ARCHAEOLOGY

When building a network model for archaeological problems, the first decision is to define what the vertices in our network are going to represent. For archaeology, one of the most natural ways to use a network picture is to represent each physical site by a unique vertex in the network. For simplicity, I will focus only on the geographical location of such sites. Alternative representations include those where a vertex represents one artefact or one

collection of artefacts, such as one type of artefact or those artefacts found at one physical site. Many of the comments and methods discussed here using geographical site locations can be translated into such contexts.

The next stage is to define the edges in our representation, so the question becomes: given a set of sites and their locations, can we understand their interactions by creating a network of edges? Here we focus on relationships derived from the geographical distances between the physical sites. Again alternative types of ‘distance’ may be derived. For instance, one may use the similarities between numbers or features of artefacts found at sites (e.g., Mills et al. 2013; Sindbæk 2007; Terrell 2010; and the chapter by Peeples et al. in this volume 2016).

This choice for our vertices immediately defines a crucial measure for the relationship of every pair of vertices—geographical distance. Even in the modern world, the proposed ‘death of distance’ brought about by modern globalization has been found to be greatly exaggerated and geography continues to play an important role. Geography is then a crucial aspect of archaeological networks. So whatever other information we add to our network picture, we must first define what we mean by distance.

Different Distances

There are two main types of distance: physical distance and ranked distance. The classic example of physical distance is the length in kilometres ‘as the crow flies’ between two points—the distance given when using a ruler on a map to measure the separation on a map. For the four sites of Figure 8.1, the direct distances are given in square brackets in Table 8.1. However, while this is the simplest measure of separation, one could imagine that the actual shortest practical route in kilometres may avoid certain difficult geographical features, such as mountain ranges, and we could use the length of such routes. Another




Figure 8.1. Four Minoan sites used for examples (data taken from Knappett et al. 2011, available from Evans et al. 2013). The straight lines indicate the direct (as-the-crow-flies) routes to and from Knossos. The dotted lines show the quickest route we penalize land travel by a factor of 3. Figure taken from Evans (2013).

Table 8.1. Physical distance measures, using quickest route or, in brackets, the direct (as the crow flies) distance (data taken from Knappett et al. 2011 and available from Evans et al. 2013). For the quickest route, land travel is penalized by a factor of three (a ‘friction’ factor) to account for slow travel over land. Thus the quickest distance typically uses coastal and sea routes rather than going overland. Table taken from Evans (2013).

Effective [<i>direct</i>] distance	Knossos	Gournia	Chania	Akrotiri
Knossos		100 km [62 km]	130 km [107 km]	129 km [119 km]
Gournia	100 km [62 km]		196 km [168 km]	142 km [143 km]
Chania	130 km [107 km]	196 km [168 km]		155 km [156 km]
Akrotiri	129 km [119 km]	142 km [143 km]	155 km [156 km]	

upgrade is to measure the ‘distance’ in terms of the time taken. For many ancient situations, travel by land was considerably slower than by water, so a sea route could involve more kilometres but actually be faster and hence a true measure of the separation. An estimate of such a quickest travel time is given in Table 8.1 for the four sites of Figure 8.1. Such estimates for the fastest walking route taking into account the extra time involved when there are significant gradients can be obtained using geographical information systems (GIS) software. Note that such distance measures need not be symmetric: walking downhill is faster than going uphill, travelling along a river upstream is slower than moving downstream. Perhaps the ultimate measure would be to estimate the cost of each route, allowing all factors to be accounted for: those noted above which could also encode issues such as the capital cost of transport, the relative safety of different routes, the local taxes imposed, and so forth. For my purposes I will leave these issues aside and I will assume that the distance between sites has been specified.

Once we have defined the physical distance between sites, we can still decide to define our model in an alternative manner, that of ‘ranked distance’. In the simplest version, taking each site in turn as the source (the origin) of a trip, we assign its nearest neighbour as the target (the destination) to be distance one, the second nearest neighbour to be distance two, and so forth. The idea is that in some circumstances we will always make a journey, irrespective of the actual distance. For instance, we may prefer to visit the nearest hospital in an emergency, however far away it is, or perhaps we have the capacity to deal with only three separate trade routes and so we choose the three closest.

A more sophisticated version of a ranked distance is the intervening opportunities distance measure, used by Stouffer (1940). This takes account of the ‘size’ of sites, defined in any appropriate manner. In this case the distance from a source site, s , to a target site, t , is defined to be the sum of all site sizes lying between the two. For an example see Table 8.2.

Table 8.2. The ranked distances and in brackets the intervening opportunities distance for the sites of Figure 8.1. These are calculated from the quickest route distances of Table 8.1. For the intervening opportunities calculation, each site has been given an exemplary size indicated in brackets after the name of the site, e.g. Knossos is size 3. Note that neither distance measure gives a symmetric table here or in general. Table taken from Evans (2013).

Rank (intervening opportunities) distance	-> Knossos (3)	-> Gournia (2)	-> Chania (2)	-> Akrotiri (1)
Knossos (3) ->		1 (3)	3 (6)	2 (5)
Gournia (2) ->	1 (2)		3 (6)	2 (5)
Chania (2) ->	1 (2)	3 (6)		2 (5)
Akrotiri (1) ->	1 (1)	2 (4)	3 (6)	

THEORETICAL COMPARISON OF MODELS

Once we have the set of sites we wish to work with and the distances between these sites, we need to add the edges to create a network. This is done with what I call an ‘edge model’.¹ There are many such models and the difficulty lies in knowing which one to choose for a given research problem. In this section we will look at the differences and similarities of models within a theoretical framework.

The first way to distinguish the models is to look at the different types of input and output they produce. These are summarized in Table 8.3 which shows considerable differences between models but also many shared features.

One of the biggest differences is in the way physical or ranked distances are chosen for each model. The discussion above made it clear that we could easily exchange one for another in any model and this reveals some simple relationships. In maximum distance networks (MDNs), a source site connects to all others up to a maximum value, say 110 km (see Figure 8.2). If we do that in terms of ranked distance, we might say connect to everyone up to ranked distance 1: the proximal point analysis (PPA) $k=1$ model, that is to your nearest neighbours (see Figure 8.2). However in the MDN and PPA models, edges are set as a function of distance through what is called a ‘deterrence function’ (discussed more generally below after Eq. 8.1). What we are doing in both cases is using a deterrence function which is a simple cut-off; the deterrence function is 1 if the distance is below a specified threshold, and 0 (indicating no edge) otherwise. The MDN and PPA models differ only in the type of distance measure, physical versus ranked. It is not surprising that these

¹ Edge models are distinct from ‘spatial influence’ models. The latter assign sites or regions of space to control or influence another site, usually creating a two-level hierarchy. Examples of spatial influence models would include Thiessen (Voronoi) polygons (Renfrew 1975), and the XTENT model (Bevan 2010; Renfrew and Level 1979).



Figure 8.2. Two simple networks for the four sites given in Figure 8.1, both derived using the quickest distances given in Table 8.1. On the left (a) is the MDN network for a 110 km cutoff, that is, an edge is placed if the separation is less than 110 km. On the right (b) is the PPA $k=1$ network with just one edge sourced per site, that is, an edge is placed between a site and its nearest neighbour. Figure taken from Evans (2013).

two models produce the same type of networks, that is simple networks where edges are present or not, and have no directions or weights.² Table 8.3 helps us see the relationships between models but we can be more quantitative in our theoretical comparison.

Functional Flow Form

Many edge models can be specified by giving a general formula for the flow between any two sites. That is, I will treat a directed edge from a source site s to target site t as representing a ‘flow’ F_{st} equal to the weight (value) of the edge, with no edge equivalent to a zero flow value, $F_{st}=0$. What this flow represents (goods, number of trips, measures of social exchange) depends on the context and choice made when modelling and are not relevant to the generic discussion here. Should the edges not be directed, then I simply set the flow to be equal in both directions $F_{st}=F_{ts}$. If the edges carry no explicit value, it is sufficient to treat the weights and flows as equal to 1. Note that in this latter case, the edges will typically represent the existence of a strong relationship rather than quantifying an actual flow. I will use the language of flows as many of these models are described elsewhere in the context of modern transport using this language. In network theory, this flow matrix is just the adjacency matrix for our network.

In many cases the ‘flows’ defined by our edge model can be made to fit into relatively simple generic form. Namely the flow from source site s to target site t is F_{st} which is given as:

$$F_{st} = a_s b_t S_s T_t f_{st} \quad (\text{Eq. 8.1})$$

² For both models, it is assumed that we also ignore the direction of the link. That is, if the distance measure is not symmetric, and it rarely is for ranked distances, then for some pairs of sites the distance may be above the threshold in one direction, and below the threshold in the other. Keeping these directions corresponds to defining new edge models.

Table 8.3. A selection of edge models which give spatial networks. If the vertex size is given as ‘fixed’ then the sizes of sites matter and can be unequal, whereas ‘equal’ indicates that the site size plays no role. A ‘variable’ site size indicates that a direct output of the model is a site size variable. These are to be distinguished from the secondary measures of site ‘size’ or at least of importance, which can be given *posthoc* in any model. References are exemplary and more references along with more details on many of the models can be found in Evans et al. 2012. There are many other examples of these models in the literature along with many other variations and yet more models, so this list provides a tiny sample from the literature.

Model	Reference	Distance Type	Site input/output Constraints	Site size S_s, T_t	Deterrence function $f(x)$	Network type
Maximum distance network (MDN)	Evans et al. (2012)	Physical d_{st}	None $a_s=1, b_t=1$	Equal	Threshold $\theta(g - x)$	Simple
Simple gravity models	Wilson, (1967)	Physical d_{st}	None $a_s=1, b_t=1$	Fixed	Any Decreasing	Weighted Directed Dense
Doubly constrained gravity model	Wilson (1967)	Physical d_{st}	Input & output a_s, b_t non trivial	Fixed	Any Decreasing	Weighted Directed Dense
Rihll & Wilson gravity models	Rihll & Wilson (1987, 1991)	Physical d_{st}	Input & output a_s, b_t non trivial	Variable	Any Decreasing	Weighted Directed Dense
Alonso models	Alonso, (1978)	Physical d_{st}	Input &/or output a_s, b_t non trivial	Variable	Any Decreasing	Weighted Directed Dense
Proximal point analysis (PPA)	Terell (1976)	Simple ranked D_{st}	Minimum input/ output	Equal	Threshold $J(g - x)$	Simple
Radiation model	Simini et al. (2012)	Intervening opp. Distance D_{st}	Output only $a_s=1,$ b_t non trivial	Fixed	Power law $\frac{1}{(D_{st} - T_t)D_{st}}$	Weighted Directed Dense

(continued)

Table 8.3. Continued

Model	Reference	Distance Type	Site input/output Constraints	Site size S_s, T_t	Deterrence function $f(x)$	Network type
Intervening opportunities model	Stouffer (1940)	Intervening opp. Distance D_{st}	Output only $a_s=1$, b_t non trivial	Fixed	Any Decreasing	Weighted Directed Dense
ariadne	Knappett et al. (2008, 2011)	Physical d_{st}	Output only $a_s=1$, b_t non trivial	Variable	$\frac{1}{(1+(d_{st}/D)^4)^2}$	Weighted Directed Dense




Figure 8.3. Illustration of the flow F_{st} from site s to site t (black arrow) in the form given by equation 8.1. The total flow into site s , the sum of all arrows pointing to s , may be specified by the input parameter S_s . The total flow out of a site t , the sum of all arrows pointing away from t , may be specified by the input parameter T_t . To enforce a constraint on outputs a parameter a_s is needed for each site but these are specified by the model through equation 8.2 and not from data. Likewise to enforce the constraint on inputs, each site must have an internal model parameter b_t set by equation 8.3. Figure taken from Evans (2013).

This is illustrated in Figure 8.3. Here the number of trips leaving the source site is S_s and the number of trips arriving at the target is T_t . These are input parameters specified by the modeller and in many cases they are chosen to be equal so that each site has a single parameter $S_s = T_s$. These parameters are some measure of the ‘size’ of each site. It is another choice in our modelling whether we set these S_s and T_t parameters equal to any measure of site size we may have (such as an estimate of the area occupied by a site) or if we choose a more complicated relationship between these parameters and data. For modern data, we may not know, for example, the number of workers and the number of jobs at each location, our S_s and T_t parameters, but we could call upon some theory to link population size or other measurable quantities, such as the length of a road, to our input parameters in a non-trivial way (see Bettencourt et al. 2007). I will focus on the generic aspects of the representation rather than such links to data.

The a_s and b_t parameters are defined by the model and are not specified directly by the modeller. They are used to impose constraints on the input and/or outputs from each site. Take the parameter a_s . In more sophisticated models, we may wish to impose the constraint that the output from each site—the total flow leaving a source site s —is guaranteed to be equal to the site size parameter S_s , which is the number of trips leaving site s . That is we demand $S_s = \sum_t F_{st}$ and it can be shown that this constraint then sets the value of a_s to be given by the formula:

$$\frac{1}{a_s} = \sum_t b_t T_t f_{st} \quad (\text{Eq. 8.2})$$

If no such constraint on outputs is imposed in the model then the a_s parameter is set to 1 and would typically not be written explicitly.




Figure 8.4. Illustration of typical forms for the deterrence function f_{st} of equation 8.1 as a function of the distance between two sites. The x axis is the distance between the sites relative to some standard scale. The deterrence functions are all 1 at short distances and decrease to 0 at large distances. Illustrated are the exponential (solid line), power law with slower fall off (dashed line, $1/(x+x_0)^2$) and a threshold or sharp cut off (dotted line). Here parameters are chosen to ensure the curves cross at $(1.0, 0.5)$. Figure taken from Evans (2013).

Any constraint on the inputs to each site is imposed in a similar way through the parameter b_t . If we demand that $T_t = \sum_s F_{st}$ is the number of trips arriving at site t then we must have:

$$\frac{1}{b_t} = \sum_s a_s S_s f_{st} \quad (\text{Eq. 8.3})$$

Again, if no such input constraint is imposed then the parameter b_t is not present (i.e. it is set to 1).

Finally, the values f_{st} are known as the deterrence function. In practice, this is usually given as a simple formula $f(d_{st})$ which depends on the distance d_{st} between the source and target site, s and t respectively. The idea is that interactions between sites should become less frequent the further apart the sites are, so these functions should decrease with distance. Typical shapes used

in the literature are shown in Figure 8.4. For instance, the MDN and PPA models mentioned earlier use the threshold form.

This deterrence function is imposing the constraint of ‘cost’, though this is in some unspecified measure of cost, rarely an actual monetary value. Typically, we imagine that the system as a whole will try to minimize its total cost while maintaining the maximum benefits. A low value for the deterrence function corresponds to a high cost and this is why we typically want the deterrence function to become smaller as the distance gets larger.

At the simplest level, by writing the different models in terms of flows of the generic form in equation 8.1, we can make a basic theoretical comparison of various edge models as shown in Table 8.3. In fact, edge models of this form can all be placed within a deeper theoretical framework known as maximum entropy, as first noted by Wilson (1967). It is a core principle of statistical physics and thermodynamics that entropy always increases in the physical world; chapter 2 of Ball (2004) has a nice non-technical description of this phenomenon. If we choose to define the ‘cost’ of a solution in a model to be minus the entropy of that particular solution, then our original viewpoint of trying to minimize some cost matches this fundamental principle of entropy maximization. However, while entropy can provide a useful view point and an alternative mathematical prescription for each model, I will not discuss this approach in more detail here (see Evans et al. (2012) for such details).

COMPARING NETWORKS NUMERICALLY

We can see just on theoretical grounds that there are a wide range of different edge models. So we can have the same set of sites and the same set of physical distances between the sites, yet each different edge model will give us a different network. How do we know which model is ‘best’? How may we compare the different models? Are the differences just different types of representation of networks which function in the same way? After all, they all have the same input. Which models are telling me something different and why?

In modern cases, we may have some reliable data on the flow between each pair of sites and a large number of such flows. The data given in a modern census on the normal commuting practices of a population is a good example of a modern dataset which can be used to test edge models (Bamis 2012; Simini et al. 2012). The basic idea would be to take the actual flow between each pair of sites and to compare that against the model’s prediction for that edge. In principle, there are standard statistical tests which measure how good this fit is. However, doing the comparison edge by edge in practice is rather more difficult (Bamis 2012).

Site Measures

Here I follow a different route and focus on the sites. This is because we have very little information on the size of flows along edges in most archaeological cases. My main objective is to compare two different models for the relationships, rather than a model and some data. Since the types of network can be so different, from simple sparse networks to dense weighted networks, comparing the edges directly (e.g. as in Bamis (2012)) will not work. Though models may look different, they may actually function in the same way, given they used the same input. So it makes sense to focus on how the network functions. To do this, I study the properties of the sites in each model.

The approach I suggest is to start by measuring quantities associated with vertices which reflect the way realistic processes react to the network. Shortest paths are a classic example of how pairs of edges are related through the network. We might then use the betweenness of a vertex (the number of shortest paths passing through a vertex) and argue that two networks are capturing the same information if the betweenness of every vertex is similar. The problem with a measure such as betweenness is that there is no generally accepted or commonly used version for weighted networks, so we could not calculate it for all edge models.

So we need to turn to measurements which can be defined on all types of networks, directed or undirected, weighted or not. There are many such measures and any could be used as the basis for network comparison of the type we describe. In practice, the choice will reflect the problem at hand and needs to be robust in the sense described by Marten Düring (2016) in Chapter 5. For simplicity and for definiteness, we will focus on two such measures, PageRank and what I call ‘influence betweenness’. I use these because they illustrate some general principles and ideas behind many of the most useful measures. The common element of the two suggested measures is that of random walkers, a concept often used in statistical physics. For this we imagine a process where a fictitious actor, the random walker, is placed on a node. At each tick of an imaginary clock, the random walker chooses to follow one edge where the choice is made in proportion to the weight of outgoing edges. We then imagine repeating these experiments many times and we look at the properties of such walkers on average. The simplest measure is to ask how often such random walkers pass through a site. This count is the PageRank measure of the ‘importance’ of that vertex. This was the basis for Google’s ranking of web pages (Brin and Page 1998), that is, the higher the PageRank of a page, the higher up the list it appeared when searching on Google.

We do not imagine that such random walkers are a realistic model of how a network is used. However, as Google’s success shows, such fictitious and unrealistic probes of the network can give useful statistical information. The

behaviour of these imaginary walkers is sensitive to the whole structure of the network without having any knowledge of this structure. While real human actors may probe a network in a less predictable manner, they also generally have poor knowledge of the whole network just like the random walkers. Real actors have different expertise, different knowledge gaps, and different biases. So perhaps on longer time scales or when looking at averages over many actors, the statistical properties of completely random walkers may not be so different from real actors, given the differing and incomplete knowledge of real actors.

The second measure I use in my examples is influence betweenness. This again uses random walkers but now imagine that at each step the walker either dies or, with equal likelihood, follows an edge chosen at random. This means that on average, these ‘influence’ walkers make three steps. Because of their finite range, such random walkers probe the ‘influence’ a site has on its neighbours as spread through the relationships defined by the network. I also set this up so that the number of walkers starting from each site is proportional to the site weight. Then, like the usual betweenness defined in terms of shortest paths, the influence betweenness is defined to be the number of these short-range random walkers passing through a site.

Similarity of Site Properties

The aim is to produce a measure of similarity for two networks based on similarity of the properties of their sites. For each network, we list the values of each measure at each site in the same order, and we can treat these lists as a vector of values representing the behaviour of the network (see Table 8.4). A similarity measure is simply one which is a maximum when the two vectors

Table 8.4. The site vectors for the four sites in Figure 8.1 (K=Knossos, G=Gournia, C=Chania, A=Akrotiri), for PR (PageRank) and IB (influence betweenness) measurements (each normalized to sum to 1). The two models are the MDN 100km and PPA k=1 models shown in Figure 8.2. The last two columns are the mean and standard deviations for the eight values for that model. The last row gives the contribution to the Pearson correlation from the value in column i , $(x_i - \langle x \rangle)(y_i - \langle y \rangle)/\sigma_x\sigma_y$ where x is the vector of eight values in the first row, y the same for the second row. Table taken from Evans (2013).

Model	K.PR	G.PR	C.PR	A.PR	K.IB	G.IB	C.IB	A.IB	MEAN	S
MDN 100 km	0.43	0.43	0.07	0.07	0.5	0.5	0	0	0.25	0.23
PPA k=1	0.5	0.17	0.17	0.17	0.61	0.13	0.13	0.13	0.25	0.19
Pearson	1.01	-0.33	0.33	0.33	2.02	-0.68	0.68	0.68	0.51	

are identical and minimum when the vectors are as different as possible. We will conjecture that the similarity of the vectors tells us about the similarity of the networks used to derive the vectors. There are many similarity measures and for a given problem one might be preferred over another. However for this generic discussion, I will use one of the most popular: the Pearson correlation coefficient. I use this as it is often a good measure: it is well known and understood, it is easily computed in computer packages, and so it provides a simple illustration of this approach. If we have two vectors, x and y , of equal length then the Pearson correlation coefficient for these two vectors is:

$$\rho_{xy} = \sum_{\text{sites}} \frac{(x_i - \langle x \rangle)(y_i - \langle y \rangle)}{\sigma_x \sigma_y} \quad (\text{Eq. 8.4})$$

where $\langle x \rangle$ and $\langle y \rangle$ are the average values for the two vectors. The Pearson correlation is a similarity measure with a maximum value of $\rho=+1$ for networks which are identical while for networks whose sites have opposite values for the chosen measurements in the vectors, these have lowest similarity value $\rho=-1$.

Site Data

Having established how to compare models through the properties of their sites, we still need to provide the basic input used by spatial models in archaeology: the locations of the sites. There are several different approaches to this aspect. If a researcher has a particular dataset for a project, it could be straightforward to do an initial comparison of models in the manner I am suggesting on the actual dataset of interest. Alternatively a standard dataset can be used, such as the thirty-nine Minoan sites used in Knappett et al. (2011) and Evans et al. (2013). However, here I choose to work with an example generated artificially (see Fig. 8.5). This approach is useful if one wants to do a systematic theoretical study.

Comparison of Networks from One Model

With the site locations specified, the methods outlined above based on site properties will now allow a comparison to be made between different models. To illustrate my approach I will start by looking at the different networks generated by just one model, the MDN model, but for different parameter values within that model. The MDN model is deterministic, that is for each value of its only parameter, its distance cut-off scale D , we get one and only one network. The model has some notable features with our test sites of




Figure 8.5. An artificial layout of forty test sites, arranged in four groups of ten. The centres of the four groups are about 150 km apart. Around each centre, ten sites are scattered randomly on a distance scale of about 50 km. The closest pair of nodes are 4km apart while the largest separation is 444 km. Figure taken from Evans (2013).

Table 8.5. Pearson correlation coefficients ρ (Eq. 8.4) for some of the MDN networks produced by the test sites of Figure 8.5. The rows and columns indicate the cut-off scale used for a particular MDN network. Table taken from Evans (2013).

r	60 km (5)	70 km (6)	80 km (7)	90 km (8)	100 km (9)
60 km (5)	1	0.93	0.88	0.87	0.87
70 km (6)	0.93	1	0.93	0.9	0.89
80 km (7)	0.88	0.93	1	0.95	0.93
90 km (8)	0.87	0.9	0.95	1	0.98
100 km (9)	0.87	0.89	0.93	0.98	1

Figure 8.5. We first get a complete loop (one which follows the boundary of the underlying square) at $D=82$ km, though the size of the largest component is rising quickly from $D=70$. Looking at the results for our similarity measure, the Pearson correlation values, we can see a few features in a simple table. For instance, with the rows and columns of the table arranged in order of the cut-off scale used in the MDN model, as in Table 8.5, we see the values generally fall away from 1 as you move away from the leading diagonal. That is, when we compare two MDN models on the same site location data but with different cut-off distance scales, the larger the difference in the cut-off distance, the lower the Pearson coefficient ρ , and so the less similar two MDN networks are. This is, of course, not surprising as changing the cut-off distance by 10 km in the MDN model for the sites of Figure 8.5 will add new edges. So to a first approximation, we expect a noticeable change in the network for each 10 km increment in cut-off distance.

However, for more general problems reading results off such a table is not so easy. We can have more than one parameter so we cannot order our models and

table in a simple fashion as there is no natural order. Luckily, there are many standard ways to extract information from large similarity tables such as I have here. All these methods define groups of networks which are similar. Here, we would expect most pairs of network in a group to have a Pearson correlation value close to 1.0. This gives us a classification of the distinct types of network which can arise for our problem, which is our primary goal. Let us illustrate two common ways to find and visualize such groups given a similarity table.

First, let us apply what is known as principal component analysis (PCA). For our visualizations, it is sufficient to know that PCA gives each network two coordinates enabling us to represent it as a point on a simple plot. These are two coordinates that describe where the network should be placed in a space of all networks.³ The coordinates are chosen such that networks which are similar should be placed close together in terms of this artificial two-dimensional space, and those which are distinct are placed far apart. This is not guaranteed but in practice it is often helpful. The other thing to bear in mind is that the first coordinate, the horizontal axis, represents the greatest variation between networks. In Figures 8.6 and 8.8 showing the PCA results, the horizontal coordinate generally represents over 60 per cent of the difference between networks while the vertical axis is less significant, representing about 25 per cent or less of the difference between two networks. So in Figures 8.6 and 8.8, height differences between two networks are less significant than horizontal differences.

Armed with this knowledge, we may look at the plot in Figure 8.6 for fourteen MDN models with the distance parameter D varying from 20 km to 150 km in 10 km increments. What we see is that the PCA visualization does indeed spread these networks out, and places them in a line of increasing threshold parameter D . Thus the method distinguishes MDN networks derived from the same data but using different thresholds. That is exactly as expected: we do not expect the same model at different parameter values to be similar. However, even though this MDN example was used for exemplary purposes and does not address our primary question, we still get useful information from the PCA visualization of Figure 8.6. The most noticeable aspect is the way the networks of large D cluster close together, especially in terms of their horizontal separation. All the networks with $D=90\text{km}$ and above are placed close together, indicating that they are very similar. In this simple example, we can understand what is happening. Once all points are well connected locally, the properties of the networks do not change very

³ There are as many coordinates as there are networks but the methods choose the two coordinates which show the largest difference between networks. We could work with more coordinates but two leads to simple visualizations. As we are projecting a large number of dimensions onto two, we cannot be sure the 2D picture always captures the true similarities and differences well, but it often works in practice.




Figure 8.6. A comparison of the MDN model on the artificial four groups of 10 sites of Figure 8.5. The MDN distance parameter D is increased from 20 km to 150 km in steps of 10 km, with models numbered from 1 to 14 in order of ascending D parameter. The plot shows the results of PCA. In this case around 91% of the differences between models are captured in just two coordinates. Also differences in the horizontal coordinate represent much larger differences between networks than similar differences in height. These results are derived from the matrix of Pearson correlation measures ρ of the vector of PageRank and influence betweenness for the forty test sites of Figure 8.5. It is clear that once connected ($D > 82$ km, labels 8 upwards), there are much smaller differences between networks than when they are disconnected. Figure taken from Evans (2013).

much. Extra links are somewhat redundant. The fact that we get a complete loop round the imaginary square (formed by the centres of the four site groups in Figure 8.5) at $D=82$ is an indicator of exactly that, as at $D=82$, the last pair of vertices which are at extremes of the global structure are connected, meaning that at a local level every vertex has easy access to travel in either way around our loop. This information could be used in a number of ways in an archaeological context. Had Figure 8.5 been a real set of sites, the results would tell us that 90km is the key distance scale for regional connectivity. Any changes in the length of journeys below 90km will make a big difference to the regional network. Changes in the length of journeys above 90km make much less difference to the way the global network functions. This distance scale can then be compared against other information on typical distances covered in trips for a given historical context. For instance, in the Early Bronze Age Aegean context (Broodbank 2000) rowing technology suggests a 10km scale is relevant, while sailing only during daylight hours might suggest a 100km scale for the Middle Bronze Age Aegean problems (Knappett et al. 2008, 2011).




Figure 8.7. A comparison of the MDN model on the artificial four groups of 10 sites of Figure 8.5. The MDN distance parameter D is increased from 20km to 150km in steps of 10km, with models numbered from 1 to 14 in order of ascending D parameter. The horizontal distance scale is the distance between networks, δ defined using equation 8.5 for $\epsilon = 10^{-4}$ and using ρ Pearson correlation measures of the vector of PageRank and influence betweenness, as illustrated in Table 8.4. To read these diagrams, draw a vertical line at network separation δ and look to see how this breaks the tree above the vertical line into clusters. Each member of the cluster is guaranteed to be within the distance δ of other members where δ is the height of the vertical line. There are several ways to define the network distance between groups of networks. The details of the agglomeration method used will not affect significant features so we should look for features which persist across different definitions of a group of networks. Here two approaches are used. On the left (a), a line at distance $\delta=0.02$ gives 7 single pieces and one large group of networks (numbers 8 to 14). In this case, the network distances between *any* two networks in separate groups is always greater than 0.02. On the right (b), a distance of $\delta=0.05$ gives 6 single pieces, one cluster of two networks 6 and 7, and one large cluster of networks 8 to 14. For this method, every network in a group is within a distance 0.05 of *all* other members of the same cluster. Figure taken from Evans (2013).

Another way to find such clusters of similar networks from a matrix of correlation values is to visualize them using dendrograms as follows. We start by converting the matrix of similarity values, ρ_{mn} for networks n and m , into a value which looks more like a measure of distance δ_{mn} between networks n and m . That is, the distance is close to zero if two networks are similar, ρ close to 1.0, and this distance measure gets larger as networks become less and less similar, ρ moves down from 1.0. This is not a physical distance, but merely a number which is zero when networks n and m are identical, while for networks with no similarity the distance between them should be extremely large. The distance should rise smoothly as the similarity of networks drops.

One simple way to produce a suitable measure of distance from the Pearson correlation values is to use a ‘network separation’ δ where:

$$\delta_{nm} = \frac{2 + \epsilon}{\rho_{nm} + 1 + \epsilon} \quad (\text{Eq. 8.5})$$

Here ϵ is a small positive number which sets the largest distance, the distance given to the most distinct networks which is roughly $(\rho_{\text{identical}} - \rho_{\text{distinct}})/\epsilon$ (about $(2/\epsilon)$ for the Pearson case). Typically we would choose ϵ to be between zero and the lowest non-zero value of $(1+\rho_{mn})$.

This gives us a matrix of distances between networks, distances measured in some theoretical space of networks. We can now create clusters by grouping networks together if the distances between them satisfy some criteria. Two such methods are illustrated in Figure 8.7. While different clustering methods give slightly different results in detail, the important results are the ones that emerge whichever method we use. Here, we see again that the networks connected with large D (number 8 to 14) are always the closest in terms of our measure of network distance δ and thus most similar. This is quickly read off the dendograms shown in Figure 8.7 and is consistent with our previous conclusions from PCA.

Comparison of Different Edge Models

We are now ready to start comparing different edge models. However, we soon encounter a difficult problem. Every model can be seen as part of a family of models (*pace* Simini et al. 2012) and so we have to choose some parameters, explicit or otherwise, to specify which models we look at. As we have seen with MDN, changing parameters within the model is not completely trivial: some changes produce big differences in the performance of the networks, some do not. Such changes can be a useful way to mimic changes seen in the historical record, for instance changing D might enable us to study how changes in technology, and hence in the range of possible interactions, could have changed the type of interactions possible. The introduction of sail technology is one example of this.

However, the problem comes when we want to compare distinct models and we need to match the input parameters and other modelling choices made for each model. Even where there appears to be a natural relationship between a physical quantity and a model parameter, such as the distance scales often found in models, we have seen with MDN that such relationships are not simple equalities and so there is no guarantee that a 100km distance scale in one model corresponds to the same physical distance produced by setting another model’s parameter scale to be 100km. Even worse, many parameters

have no natural match to physical parameters. For instance, is a $k=4$ PPA to be compared to an MDN network with $D=50$ km, 100 km, or 150 km?

The answer again is to look for a physical measure, and to match networks with the same value of this parameter. Good choices are those which are well defined on all types of network and again, measures based on random walkers are useful here. As we are interested in spatial models here the obvious measure to look for is one of the physical distance scale captured by every network.

The idea is to measure the average physical distance a random walker travels when making a single step. Consider a single source vertex s . Then the average distance from s to a neighbouring target site t that a random walker travels is Δ_s where:

$$\Delta_s = \frac{\sum_t w_{st} d_{st}}{\sum_t w_{st}} \quad (\text{Eq. 8.6})$$

where d_{st} is the distance from s to t and w_{st} is the weight we give to an edge, usually the flow F_{st} along this edge. This is illustrated in Table 8.6.




Figure 8.8. Scatter plot of the first two components of the vertex vectors for seven models using the test data sites of Figure 8.5. These networks all have an average weighted neighbour distance of about $\Delta=57$ km. They are numbered as follows: 1=PPA ($k=4$), 2=directed PPA ($k=5$), 3=MDN ($D=82$ km), 4=simple gravity model (distance scale $D=75$ km), 5=doubly constrained gravity model ($D=45$ km), 6=Rihll and Wilson gravity model ($D=55$ km, $\alpha=1.03$), 7=Rihll and Wilson gravity model ($D=50$ km, $\alpha=1.20$). For abbreviations and references see Table 8.3. The number of edges kept in a PPA model is given by k . The directed PPA model is a PPA model where directions are retained on edges. The gravity models all use a deterrence function of the form $f(x)=(1+x^4)^{-2}$ as used in Knappett et al. (2008). Figure taken from Evans (2013).

Table 8.6. Illustration of average weighted neighbour distance for Knossos in the four site example of Figure 8.1 and using the effective distance of the quickest routes given in Table 8.1. Here 90% of the edges weight from Knossos is in the link to Gournia so the distance to this site dominates the calculation of the weighted distance to its neighbours. Table taken from Evans (2013).

Knossos to	Gournia	Chania	Akrotiri	Overall
Quickest effective distance	100 km	130 km	129 km	Average 120 km
Edge weight	0.53	0.05	0.01	Total 0.59
Weighted distance	89.8 km	11.0 km	2.2 km	Total 103 km
Contribution				

To give a physical characteristic distance for a network we just take the average over all sites, that is $\Delta = (1/N) \sum_s \Delta_s$ where N is the number of vertices in the network. This is a measure of the actual physical distance between sites as ‘seen’ by a random walker on the network, regardless of how the network is defined. It is then fair to compare networks with similar values of this average distance Δ , varying the parameters in each model to find values with a similar value of Δ . For instance in Figure 8.8, six models are studied. For one model with two parameters (Rhill and Wilson 1987, 1991) we look at two very different parameter combinations which give the same value for Δ . We see the same pattern when using the dendrogram visualizations as shown in Figure 8.9.




Figure 8.9. The similarity between seven edge models (numbered as in Figure 8.8) producing the same average weighted distance of $\Delta=57$ km. These hierarchical plots are defined in exactly the same way as Figure 8.7 using network distance δ of equation 8.5 on the vertical axis. In both cases, the networks produced from the Rihll and Wilson model (Rihll and Wilson 1987, 1991) but for different parameters seem to be outliers, producing networks which are very distinctive from the other networks and even between themselves. Figure taken from Evans (2013).

CONCLUSIONS

I have looked at different ways to compare the many models used to create spatial networks based on the locations of a fixed number of sites. Theoretical comparisons can be made comparing features of models as in Table 8.3. A more precise theoretical comparison can be made using the maximum entropy formalism, as most models can be derived from some precise form for the entropy function (or equivalently a cost–benefit function). However, I have also suggested a numerical recipe for a quantitative comparison of spatial network models which focuses on the way the networks function. The advantage of this last approach is that even models which produce very different networks could have similar functionality. After all, all the networks are based on the same geographical input. The approach is as follows:

- Choose model parameters so that all networks have the same average weighted distance, Δ of equation 8.6.
- Calculate vertex measures that probe the way the network functions, for instance using random-walker-based measures such as PageRank.
- Define a similarity matrix for networks based on the vertex values, such as the Pearson correlation matrix ρ of equation 8.4. If needed convert this into a network separation through a distance function such as δ of equation 8.5.
- Use multivariate analysis techniques such as principal component analysis (Figures 8.6 and 8.8) or hierarchical clustering methods (Figures 8.7 and 8.9) to produce conclusions and visualizations.

An advantage of the approach outlined here is that it draws on largely standard quantities and so can be implemented quickly in many numerical frameworks.

A systematic survey of the wide range of models and their parameters along with different possible spatial layouts of sites has yet to be performed. However, preliminary results using the numerical approach suggest a couple of conjectures. Firstly, the Rihll and Wilson models (1987, 1991) are often quite distinct from other models. Indeed, the original use of this model was as a zone-of-control type model such as Thiessen polygons (Renfrew 1975) and the XTENT model (Renfrew and Level 1979). These models are not trying to describe regional interactions, just local dominance. A second hypothesis is that one of the most important features of any model is the constraints on the inputs and/or outputs of each site. A corollary is that the type of distance measure used, a physical distance versus a ranked distance, seems to be a much less important choice in most practical situations.

The archaeological payoff of this work will be that the framework allows us to understand the similarities and differences between the wide range of

network models available. Once this information has been gathered, researchers will be able to select the best model for their task in a systematic manner. In the future, each family of models defined by the methods outlined here could be linked to particular types of problem, simplifying model selection. Further, we can test the reliability of our conclusions against modelling uncertainties as any results which appear in all models from the same family are likely to be reliable. In this way, the methods given here help address the real concerns that researchers have when trying to extract robust results from archaeological models, adding to the approaches in the chapters by Düring (2016), Peeples et al. (2016), and Rivers (2016) in this volume.

ACKNOWLEDGEMENTS

This work follows on from an initial study of this topic with Giovanni Brandani (Brandani 2010). I would like to thank Elias Bamis, Giovanni Brandani, Michael Gastner, Carl Knappett, and Ray Rivers for many conversations on different aspects of these problems.

REFERENCES

- Ball, P. (2004). *Critical Mass: How One Thing Leads to Another*. London: William Heinemann.
- Bamis, I. (2012). ‘Constrained Gravity Models for Network Flows’, MSc thesis (Imperial College London).
- Bettencourt, L. M. A., Lobo, J., Helbing, D., Kühnert, C., and West, G. B. (2007). ‘Growth, innovation, scaling, and the pace of life in cities’, *Proceedings of the National Academy of Sciences*, 104:7301–6.
- Bevan, A. (2010). ‘Political geography and palatial Crete’, *Journal of Mediterranean Archaeology*, 23:27–54.
- Brandani, G. B. (2010). ‘Complex Networks: Rihll and Wilson Model and Archaeology’, BSc Project thesis (Imperial College London).
- Brin, S. and Page, L. (1998). ‘The anatomy of a large-scale hypertextual Web search engine’, *Computer Networks and ISDN systems*, 30:107–17.
- Broodbank, C. (2000). *An Island Archaeology of the Early Cyclades*. Cambridge: Cambridge University Press.
- Düring, M. (2016). ‘How Reliable are Centrality Measures for Data Collected from Fragmentary and Heterogeneous Historical Sources? A Case Study’, in T. Brughmans T., A. Collar, and F. Coward (eds), *The Connected Past: Challenges to Network Studies in Archaeology and History*. Oxford: Oxford University Press, 85–101.
- Evans, T. S. (2004). ‘Complex Networks’, *Contemporary Physics*, 45:455–74.

- Evans, T. S. (2013). 'Background data for "Which Network Model Should I Use?"', [electronic repository], <http://dx.doi.org/10.6084/m9.figshare.783864> (accessed 29th July 2013).
- Evans, T. S., Knappett, C., and Rivers, R. J. (2012). 'Interactions in Space for Archaeological Models', *Advances in Complex Systems*, 15:1150009.
- Evans, T., Knappett, C., and Rivers, R. J. (2013). 'Thirty Nine Minoan Sites', [electronic repository], (accessed 8th July 2013).
- Knappett, C., Evans, T. and Rivers, R. J. (2008). 'Modelling maritime interaction in the Aegean Bronze Age', *Antiquity*, 82:1009–24.
- Knappett, C., Evans, T., and Rivers, R. J. (2011). 'The Theran eruption and Minoan palatial collapse: new interpretations gained from modelling the maritime network', *Antiquity*, 85:1008–23.
- Lane, D., Van der Leeuw, S. E., Pumain, D., and West, G. (2009). *Complexity Perspectives in Innovation and Social Change*. Dordrecht: Springer.
- Mills, B. J., Clark, J. J., Peeples, M. A., Haas, Jr, W. R., Roberts, Jr, J. M., Hill, J. B., Huntley, D. L., Borck, L., Breiger, R. L., Clauset, A., and Shackley, M. S. (2013). 'Transformation of social networks in the late pre-Hispanic US Southwest', *Proceedings of the National Academy of Sciences*, 110:5785–90.
- Peeples, M. A., Mills, B. J., Haas, W. R. J., Clark, J. J., and Roberts, J. M. J. (2016). 'Analytical Challenges for the Application of Social Network Analysis in Archaeology', in T. Brughmans, A. Collar, and F. Coward (eds), *The Connected Past: Challenges to Network Studies in Archaeology and History*. Oxford: Oxford University Press, 59–84.
- Renfrew, C., (1975). 'Trade as action at a distance', in J. Sabloff, and C. Lamberg-Karlovsky (eds), *Ancient Civilization and Trade*. Albuquerque: University of New Mexico Press, 3–59.
- Renfrew, C., and Level, E. (1979). 'Exploring dominance: predicting polities from centres', in C. Renfrew, and K. L. Cooke (eds), *Transformations: Mathematical Approaches to Culture Change*. London: Academic Press, 145–67.
- Rihll, T. E., and Wilson, A. G. (1987). 'Spatial interaction and structural models in historical analysis: some possibilities and an example', *Histoire & Mesure*, 2:5–32.
- Rihll, T. E., and Wilson, A. G. (1991). 'Modelling settlement structures in ancient Greece: new approaches to the polis', in J. Rich and A. Wallace-Hadrill (eds), *City and Country in the Ancient World*. London: Routledge, 59–95.
- Rivers, R. (2016). 'Can Archaeological Models Always Fulfil our Prejudices?', in T. Brughmans, A. Collar, and F. Coward (eds), *The Connected Past: Challenges to Network Studies in Archaeology and History*. Oxford: Oxford University Press, 123–47.
- Rivers, R., Knappett, C., and Evans, T. (2013). 'Network Models and Archaeological Spaces', in A. Bevan and M. Lake (eds), *Computational Approaches to Archaeological Spaces*. Walnut Creek, CA: Left Coast Press.
- Simini, F., Gonzalez, M. C., Maritan, A., and Barabasi, A.-L. (2012). 'A universal model for mobility and migration patterns', *Nature*, 484:96–100.
- Sindbæk, S. M. (2007). 'Networks and nodal points: the emergence of towns in early Viking Age Scandinavia', *Antiquity*, 81:119–32.
- Stouffer, S. A. (1940). 'Intervening Opportunities: A Theory Relating to Mobility and Distance', *American Sociological Review*, 5:845–67.

- Terrell, J. E. (2010). 'Language and Material Culture on the Sepik Coast of Papua New Guinea: Using Social Network Analysis to Simulate, Graph, Identify, and Analyze Social and Cultural Boundaries Between Communities', *Journal of Island and Coastal Archaeology*, 5:3–32.
- Wilson, A. G. (1967). A statistical theory of spatial distribution models', *Transport Research*, 1:253–69.

9

Networks, Homophily, and the Spread of Innovations

Anne Kandler and Fabio Caccioli

INTRODUCTION

The question of how and why innovations spread through populations has been the focus of extensive research in various scientific disciplines over recent decades. Generally, innovation diffusion is defined as the process whereby a few members of a social system initially adopt an innovation, then over time more individuals adopt until all (or most) members have adopted the new idea (e.g. Rogers 2003; Ryan and Gross 1943; Valente 1993). Anthropologists and archaeologists have argued that this process is one of the most important processes in cultural evolution (Richerson et al. 1996) and much work has been devoted to describing and analysing the temporal and spatial patterns of the spread of novel techniques and ideas from a particular source to their present distributions. Classic case studies include the spread of agricultural inventions such as hybrid corn (e.g. Griliches 1957; Ryan and Gross 1943), the spread of historic gravestone motifs in New England (Dethlefsen and Deetz 1966; Scholnick 2012), and the spread of bow and arrow technology (Bettinger and Eerkens 1999). (For a more comprehensive list see Rogers and Shoemaker (1971) who reviewed 1,500 studies of innovation diffusion.) Interestingly, the temporal diffusion dynamic in almost all case studies is characterized by an S-shaped diffusion curve describing the fraction of the population which has adopted the innovation at a certain point in time. Similarly, the spatial dynamics tend to resemble travelling wave-like patterns (see Steele 2009 for examples).

The basic puzzle posed by innovation diffusion is the observed lag between an innovation's first appearance and its general acceptance within a population (Young 2009). In other words, what are the individual-level mechanisms that give rise to the observed population-level pattern? Again, scientific fields as diverse as economics/marketing science (e.g. Bass 1969; Van den Bulte and Stremersch 2004; Young 2009), geography (e.g. Hägerstrand 1967),

or social science (e.g. Henrich 2001; Steele 2009; Valente 1996; Watts 2002) offer interesting insights into this question without reaching a consensus about the general nature of individual adoption decisions.

In archaeological and anthropological applications, population-level patterns inferred from the archaeological record, such as adoption curves, are often the only direct evidence about past cultural traditions (Shennan 2011). To explore which individual-level processes could have produced the inferred patterns researchers need to ‘reverse engineer’ conclusions about those processes in current or past populations, given the knowledge of how certain population-level patterns have changed over space and time.

In this chapter, we contribute to the debate by developing a simulation framework which explicitly models the adoption decision of individuals with regard to an innovation, and therefore establishes a causal relationship between individual-level processes and population-level patterns. Using this modelling approach enables us to (1) systematically explore the effects of, for example, different adoption rules or population structures on population-level patterns such as adoption curves, and (2) derive theoretical expectations about, for example, the shape of the adoption curve, based on the assumed adoption dynamic. Comparing those expectations with the observed data might allow us to at least exclude individual-level processes that could *not* have produced the observed data. Contrary to previous models (e.g. Henrich 2001; Kandler and Steele 2009), we relax the assumption of a well-mixed population and consider structured populations. Our framework is a generalization of the well-known Watts model (Watts 2002) where social interactions (and consequently information flow) between individuals are modelled as networks, that is, collections of nodes representing individuals and links connecting pairs of nodes representing social relations. Additionally we allow for individual heterogeneity: individuals can possess different levels of social status. Besides introducing individual heterogeneity, we also account for the possibility of a homophilistic bias, that is, the tendency for individuals to have (non-negative) ties with individuals that are similar in socially significant ways (Lazarsfeld and Merton 1954). This framework enables us to analyse how the interplay between network topology and individual heterogeneity in the form of status differences and homophily can alter the adoption dynamic of the innovation. In particular we focus on two questions:

1. Does the adoption dynamic change qualitatively depending on the social status of the innovator?
2. Does the subpopulation of high-status individuals exhibit a different adoption pattern than the subpopulation of low-status ones?

As our approach is based on the theory of networks literature, in the following we very briefly review relevant aspects of it, introduce the Watts model (Watts 2002), and summarize its main findings.

Diffusion of Innovation in Complex Networks

Much attention has been devoted to understanding how topological properties of social networks affect the diffusion dynamics of cultural innovations (e.g. novel ideas or beliefs). Valente (1996) offers a chronological overview of research into this question. The initial approach was to count the number of links each individual has with others and to correlate this variable with innovativeness (defined as the time taken for an individual to adopt the innovation) (e.g. Coleman et al. 1966; Rogers 2003). The individuals with the highest number of links were identified as opinion leaders, and were theorized to be a significant influence on the adoption dynamic (Katz and Lazarsfeld 1955; Lazarsfeld et al. 1968). However, Granovetter (1973, 1982) argued that weak ties (individuals only loosely connected to others in the network) were necessary for spread to occur across a population. Further refining the concept, Burt (1980, 1987) focused on the impact of structural equivalence (the extent to which two nodes are connected to the same others, i.e. have the same neighbourhood) on adoption dynamics, and it has been argued that network properties such as centrality, density, and reciprocity may play an important role as well (e.g. Valente 1995). For a discussion of the applicability of centrality measures to archaeological data see Chapter 5 of this volume (Düring 2016).

A different line of research considered the consequences of *thresholds*. Threshold models assume that an individual adopts an innovation based on the proportion of people in the social system which has already adopted the innovation (Granovetter 1978). Each agent i is endowed with a threshold, and it adopts if the number of individuals that have already adopted exceeds its threshold. Consequently, an individual's adoption decision is a function of the adoption decisions of others in the system. Individuals with low thresholds will adopt at a very early stage, while individuals with high thresholds will only do so if almost the whole system has already adopted. This threshold idea has been further taken forward in the seminal paper by Watts (2002) who studied the adoption dynamic of an innovation in a random network.

Watts (2002) considers a population of N individuals and assumes that each individual i must make a decision with regard to some issue X —in our case, whether to adopt an innovation or not. The N individuals are organized in a network where each individual is connected to exactly k neighbours with probability p_k and the average number of neighbours is $\langle k \rangle = z$ with z being a positive constant (in the following, z is referred to as the average degree of the network, that is the average number of neighbours per node). Systems of this type are called random networks and represent a generalization of the Erdős-Rényi network (Erdős and Rényi 1959) by allowing for general degree distributions (Newman et al. 2002).

All individuals apart from a small fraction ($\Phi_0 \ll 1$) are initialized with state 0, representing no adoption of the innovation. The fraction Φ_0 of the




Figure 9.1. Example of a population organized in a random network. The dark grey nodes represent individuals who have not adopted the innovation (initialized with state 0) whereas the light grey node represents the innovator, initialized with state 1.

population comprises all initial adopters of the innovation (innovators), which are initialized with state 1, and in the following we assume $\Phi_0=1/N$ (see Fig. 9.1 for a visualization of this initial condition). The presence of the innovation in the population perturbs the system in such a way that individuals with state 0 who are exposed to the innovation (in Fig. 9.1, the six nodes that connect with the grey, innovating node) have to decide whether to adopt the innovation (changing their state to 1) or not (so leaving their state unchanged).

In this model, the decision of an individual i to adopt the innovation depends solely on the number of individuals who have already adopted the innovation in its local neighbourhood (where *local neighbourhood* stands for the set of individuals i is connected to). Therefore i 's decision depends on the decisions of other individuals and not on its own assessment of the ‘usefulness’ of the considered innovation. It holds that: the more individuals that have already adopted, the more likely individual i is to adopt. In more detail, each individual is randomly assigned a threshold ϕ_i drawn from a distribution $f(\phi)$. This distribution is defined on the interval $[0,1]$ and normalized so that $\int_0^1 f(\phi) d\phi = 1$ but which is otherwise arbitrary (Watts 2002). The threshold ϕ_i can be interpreted as the amount of social reinforcement needed for an individual to adopt the innovation: the higher the threshold of individual i the more individuals have to have adopted the innovation before i will adopt. Consequently, individuals with a high threshold will adopt the innovation relatively late, if at all. Decision rules of this type are called threshold rules and have been derived for a variety of theoretical scenarios, including coordination

and public goods games, and adoption decision in the presence of uncertainty (Lopez-Pintado and Watts 2007). Empirical and experimental evidence is limited but tends to support the assumption that individuals apply threshold rules when making decisions in the presence of social influence (e.g. Latané and L'Herrou 1996; Young 2006). As simplistic as it may appear, this decision framework is relevant to surprisingly complex problems (Watts 2002).

To formalize mathematically the adoption dynamic described above, we denote by σ_i the state of individual i (either 0 or 1) and by $v(i)$ its local neighbourhood and define:

$$\sigma_i = \Theta\left(\frac{\sum_{j \in v(i)} \sigma_j}{|v(i)|} - \phi_i\right) \quad (\text{Eq. 9.1})$$

where $|v(i)|$ stands for the numbers of neighbours of individual i and $\Theta(x)=1$ if $x>0$ and zero otherwise. Individual i will adopt the innovation if the average number of adopters in its neighbourhood exceeds its threshold ϕ_i

$$\frac{\sum_{j \in v(i)} \sigma_j}{|v(i)|} > \phi_i.$$

The population evolves in successive time steps in which each individual updates its state in a random, asynchronous order.

Using this model set-up, Watts (2002) studied how the underlying network of interpersonal influences that drives decision-making influences the adoption dynamic of a population. Central to this question is the analysis of the size of the cascade triggered by a given initial fraction of adopters Φ_0 . ‘Cascades’ here refers to a sequence of adoption events of any size (by size we mean the final fraction of the population having adopted the innovation); a cascade is called ‘global’ if more than a fixed fraction of a large but finite network has adopted the innovation (Watts 2002). The main results of this analysis can be stated as follows (for details, see Watts 2002):

- In sufficiently sparse networks (i.e. networks with a small average number of neighbours z) cascade propagation is prevented by the low connectivity of the network.
- In contrast, if the network is sufficiently dense (i.e. it has a large average number of neighbours z) cascade propagation is prevented by the local stability of the nodes themselves (meaning $\phi_i > 1/k_i$, where k_i is the number of neighbours of individual i).
- Between these two extremes, there exists a ‘window’ of connectivity within which global cascades occur with a non-zero probability. The properties of the system are very different at the lower and upper ends of this window. Close to its lower end, the system is characterized by a




Figure 9.2. Spreading probability (dark grey dots) and final cascade size (light grey squares) in the Watts model with $N=10000$ and $\phi_i=\phi=0.16$.

powerlaw distribution of cascade size, while close to the upper end the distribution of cascade sizes becomes bimodal: either the innovation spreads through almost the whole population or the innovation does not spread at all. Figure 9.2 shows the spreading probability (dark grey dots) and the conditional cascade size (light grey squares) as a function of the degree z for a network of size $N=10000$ and $\phi_i=\phi=0.16$ for all individuals i . The spreading probability describes the probability that the innovation spreads at least through a finite fraction of the population, in our case 5 per cent. The cascade size is then calculated based on simulations in which the innovation successfully spreads through 5 per cent of the population. It is obvious that the innovation spreads with a high probability only in networks characterized by intermediate degree. We further observe that, once the innovation spreads through 5 per cent of the population, in sufficiently dense networks it will spread through almost the whole population.

- The initial adopter plays the role of an innovator and early adopters are individuals with $\phi_i > 1/k_i$. Consequently, no cascade is possible unless the innovator is connected to some early adopters. The adoption dynamic

may depend less on the characteristics of the innovator itself than on the structure of the community of early adopters.

Based on this threshold model, Watts and Dodds (2007) tested the so-called ‘influential hypothesis’ (Katz and Lazarsfeld 1955; Lazarsfeld et al. 1968) and considered the impact of influential individuals on the adoption dynamic. In this model, they interpreted the underlying random network as an influence network (the existence of a link between two individuals means that they are able to influence each other with regard to the considered issue X) and defined influential individuals as individuals with many links. Interestingly, they found that under most conditions, influential nodes are less important than generally assumed, either as innovators or as early adopters. In other words, the adoption dynamic did not change qualitatively depending on whether the innovation was introduced by a more or less influential individual respectively, or whether influential individuals were numerous in the subpopulation of early adopters.

In the following, we will add a further layer of complexity to those approaches. We assume that influence is not only manifested in the network structure but also in the adoption dynamic. The underlying network still controls the flow of information in the population (and consequently who is able to influence whom) but additionally, we allow individuals to possess different levels of social status. Status differences affect the adoption dynamics in that high-status individuals have a stronger influence on their neighbours than low-status ones (e.g. Centola 2011; Magee and Galinsky 2008). Besides introducing individual heterogeneity in the form of status differences, we are also interested in exploring the effects of homophilistic bias, that is, the tendency for individuals to have (non-negative) ties with individuals that are similar in socially significant ways (Lazarsfeld and Merton 1954). Based on the proposed status differences, we assume that individuals are more likely to interact with those of similar social status (Centola et al. 2007). We model this bias by generating networks which exhibit correlations between the statuses of neighbouring individuals. Varying the degree of these correlations allows us to systematically explore the effects of different levels of homophily on the adoption dynamics.

HETEROGENEITY IN THE POPULATION

In the following, we consider a heterogeneous population by allowing individuals to possess different levels of social status. We use the concept of social status in its most simplistic form and interpret status as a one-dimensional variable quantifying, among other things, an individual’s social and economic

attributes (see, e.g. McPherson and Smith-Lovin 1987). We assume that each individual i is assigned a status variable s_i which can take values in the interval $[0,1]$ with 0 indicating the lowest possible status and 1 the highest. Status differences are reflected in an individual's potential to influence other individuals in its neighbourhood to adopt or not to adopt the innovation. It holds: the higher the status of an individual the higher is its influence. Consequently, the adoption rule (Eq. 9.1) is changed to:

$$\sigma_i = \Theta \left(\frac{\sum_{j \in v(i)} s_j \sigma_j}{\sum_{j \in v(i)} s_j} - \phi_i \right) \quad (\text{Eq. 9.2})$$

where $\Theta(x)=1$ if $x>0$ and zero otherwise. Individual i still evaluates the social information available in its neighbourhood but now it weights its neighbours by their social status s_j . In particular, an individual i will adopt if it holds:

$$\frac{\sum_{j \in v(i)} s_j \sigma_j}{\sum_{j \in v(i)} s_j} > \phi_i$$

Equation 9.2 implies that an individual's decision does not depend on its own social status, but only on the status of its neighbours. In other words, high- and low-status individuals with the same neighbourhood behave in the same way. We note that alternative adoption rules can be easily included in the suggested framework by changing equation 9.2. For example, a more general version of the model could include a threshold $\phi_i(s_i)$ which is a function of the social-status of individual i . For the sake of simplicity, we focus in this paper on the case of a constant threshold and leave the investigation of different adoption decisions for future research.

In the following, we explore whether and how this population structure, superimposed on the network of interpersonal influences, affects the adoption dynamic of the innovation. For the sake of simplicity, we assume that individuals either possess status s_1 or s_2 with $s_1 < s_2$ and $s_1 + s_2 = 1$ and the population is therefore arranged into two subpopulations of different status. We note that the condition $s_1 + s_2 = 1$ is an arbitrary model assumption quantifying the status difference between the two subpopulations. For example, while $s_1 = s_2 = 0.5$ characterizes the situation where both subpopulations have the same status $s_1 = 2/3$ and $s_2 = 1/3$ describes the situation where subpopulation 1 has a status twice as high as subpopulation 2. Further, we assume $\phi_i = \phi = 0.16$ for all individuals i . The underlying network of interpersonal influences is modelled as an Erdős-Rényi network with $N=10000$ individuals whereby 5000 possess high status and 5000 low status. We note that the Erdős-Rényi network represents the simplest network model given the constraint of a fixed average degree. However, as this paper focuses on exploring whether and how the adoption dynamics of an

innovation is affected by individual heterogeneity, rather than by the degree distribution of the network, we consider this choice justified. The analysis presented in this chapter can be extended to networks which represent real social networks more accurately.

Adoption Dynamic

Following Watts (2002), we start by analysing the spreading probability of an innovation and the size of the resulting cascade but distinguish whether the innovation was introduced by a high- or a low-status individual. As stated previously, the spreading probability describes the probability that the innovation spreads through at least 5 per cent of the population. If the innovation is introduced by a high-status individual, Figure 9.3a shows that the spreading probabilities of the subpopulations of high-status (light grey squares) and low-status individuals (dark grey dots) exhibit a similar behaviour for varying degree z of the network. We observe that too sparse or too dense networks are characterized by low spreading probabilities. However, the reasons for this behaviour are different. As already noted by Watts (2002), in sparse networks the spread of the innovation is prevented by the low connectivity of the network, whereas in dense networks, spreading is prevented by the local stability of the nodes themselves (meaning $\phi_i > 1/k_i$, where k_i is the number of neighbours of individual i).




Figure 9.3. (a) Spreading probabilities of the high- and low-status subpopulations when i) the innovation is introduced by a high-status individual (light grey squares and dark grey dots, respectively) and ii) the innovation is introduced by a low-status individual (black triangles). We assumed $N=10000$ and $\phi_i=\phi=0.16$. (b) Corresponding average cascade size of the high-status (light grey squares) and low-status (light grey squares) subpopulations when the innovation is introduced by a high-status individual. No results are shown for the situation of a low-status innovator as no spreading occurs in this case.




Figure 9.4. Fractions of the high-status (dashed light grey line) and low-status (solid dark grey line) subpopulations that have adopted the innovation at different time steps.

Interestingly, Figure 9.3a also shows that if the innovation is introduced by a low-status individual, no or at most very little spreading occurs in either subpopulation (cf. black triangles) *regardless* of the structure of the network. This behaviour results from the fact that high-status individuals are more influential than low-status individuals. It might appear counterintuitive that the innovation does not spread within the low-status subpopulation, but in the present set-up of the model, the status of an individual does not correlate with the status of its neighbours and therefore the structure of the local neighbourhood of high- and low-status individuals is statistically identical. Furthermore, the decision to adopt the innovation (cf. Eq. 9.2) depends only on the social status of the neighbours but results in the same aggregated adoption dynamic in both subpopulations.

Figure 9.3b shows the average cascade size in the subpopulations of high-status (light grey squares) and low-status individuals (dark grey dots) if the innovation is introduced by a high-status individual for varying degree z of the network (as before, the average cascade size is calculated on the basis of the simulations where the innovation successfully spread through 5 per cent of the population). For sufficiently dense networks (those with average degrees larger than four), we observe that once the innovation spreads through 5 per cent of the population it will spread through almost the entire population.

In order to gain information about the temporal dynamics of the adoption process, we plot in Figure 9.4 the fractions of the high-status (dashed light grey

line) and low-status subpopulations (solid dark grey line) which have adopted the innovation at a certain time step t . As expected both subpopulations exhibit the same dynamics and we observe the S-shape of the adoption curves which is characteristic of almost all archaeological case studies (see, e.g. Griliches 1957; Morrill 1985).

To summarize, the adoption dynamic depends crucially on the social status of the first adopter. While no or at most very little spreading occurs if the innovation is introduced by a low-status individual, we observe different spreading probabilities depending on the density of the network if the innovation is introduced by a high-status individual. Importantly, we do not find any differences between the adoption dynamic of the high- and low-status subpopulations.

HOMOPHILY

Thus far we have based our analysis on the assumption that the underlying network specifying the interpersonal information flow in the population is random. This implies that every link between two individuals is chosen independently from, for example, the attributes of the individuals or pre-existing links. However, real social networks often show strong interdependencies deriving from different cultural mechanisms, such as social roles (Nadel 1957), group affiliations (Feld 1981), or homophily (Lazarsfeld and Merton 1954). All of those mechanisms introduce local structure into the network (Watts and Dodds 2007).

In this section, we explore the consequences of a homophilistic bias on the adoption dynamic of an innovation. Very generally, homophily is the tendency of individuals with similar traits (including physical, cultural, and attitudinal characteristics) to interact with each other more than with people with dissimilar traits (Centola et al. 2007). However, homophilistic tendencies can be expressed in different forms and in the following we consider only what is called ‘choice homophily’ (McPherson and Smith-Lovin 1987; McPherson et al. 2001), where patterns of interaction are driven by preferences for similarity. Centola et al. (2007) modelled this kind of homophily by allowing the social network to evolve as a function of cultural similarities and differences between individuals. We follow this idea and assume that individuals have a preference for interacting with individuals of the same status and therefore assume that the homophilistic bias acts on social status. We are interested in exploring how the introduction of local structure into the network influences the spread of an innovation. The adoption of the innovation has no effect on the social status of an individual.

In our model, we account for the similarity preference by assuming that individuals have a higher probability of being connected with individuals of the same status. The resulting network possesses the same degree distribution as in the previous section, but now links between individuals are arranged in such a way that individuals with the same status tend to be connected more frequently than before. Networks of this kind are called *correlated networks* and the degree of correlation can be interpreted as a measure of the strength of the homophilistic bias.

Empirically, it is often very difficult to distinguish the effects of homophily from the effects of social influence, or, in other words, to disentangle whether individuals act in similar ways due to their similarities or because they are close in the network due to some form of social influence that acts along network ties (Shalizi and Thomas 2011). However, this model is not set up to explore the origin of homophily but simply to understand the consequences of this phenomenon on the spread dynamic of an innovation.

Correlated Networks

In the following, we briefly describe a Monte Carlo algorithm to construct correlated networks. This algorithm is based on the one introduced by Noh (2007). We first generate a random network and introduce the quantity C that counts the number of links between individuals with different status:

$$C = J \sum_{i=1}^N \sum_{j \in v(i)} \delta(s_i, s_j) \quad (\text{Eq. 9.3})$$

where $\delta(s_i, s_j) = \begin{cases} 1 & \text{if } s_i \neq s_j \\ 0 & \text{if } s_i = s_j \end{cases}$ and the parameter J controls the degree of correlation in the network. Larger values of J indicates a higher correlation between the status of neighbouring nodes. In order to create a correlated network, we now apply the following rewiring algorithm:

1. Select two pairs of connected nodes (i,j) and (k,l) at random.
2. Compute the change ΔC in the quantity C that would occur when exchanging the links between the two pairs of nodes, that is, deleting the links between (i,j) and (k,l) and introducing new links between (i,l) and (k,j) .
3. If it holds $\Delta C < 0$, then exchange the links. Otherwise, exchange the links only with probability $e^{-J\Delta C}$ (It follows that for $J \rightarrow \infty$ the network is perfectly correlated as in this case $e^{-J\Delta C} \rightarrow 0$ if $\Delta C > 0$).

These rewiring steps are repeated until the network ceases to change statistically, and the resulting network shows the desired degree of correlation




Figure 9.5. Relationship between the parameter J and the network correlation.

and therefore level of homophily. Figure 9.5 illustrates how the values of the parameter J are translated into correlation levels of the network. A correlation of 0 means that individuals are linked at random, independent of their status. Contrarily, a positive (negative) correlation means that there is a tendency for individuals with the same (different) status to be connected and at the extreme (correlation of 1) both subpopulations are segregated.

In the following we explore how the results obtained in the section ‘Heterogeneity in the population’ change if we allow for status correlations in the underlying network of interpersonal influences and therefore account for a homophilistic bias in the population.

Adoption Dynamic

We begin by analysing the spreading probability of an innovation and the size of the resulting cascade in a highly correlated network ($J=2$). As before, we assume that an individual’s decision to adopt the innovation is given by equation 9.2.

Figure 9.6a shows the spreading probability of the subpopulations of high-status (light grey squares) and low-status individuals (dark grey dots) for varying degree z of the network and for the situation where the innovation is introduced by a high-status individual. Firstly, we observe that both subpopulations show similar adoption behaviour and secondly, that the network




Figure 9.6. Spreading probabilities of the high-status (light grey squares) and low-status subpopulations (dark grey dots) when the innovation is introduced by a high-status individual. We assumed $N=10000$ and $\phi_i=\phi=0.16$. (b) Corresponding average cascade size of the high-status (light grey squares) and low-status subpopulations (dark grey dots). (c and d) Same analyses as in (a) and (b) but with $N=1000$.

correlations (and consequently the homophilistic tendency in the population) change the adoption dynamics (cf. Figure 9.3 for the results in the uncorrelated situation). The spreading probabilities indicate that the spread of the innovation is only supported in a narrower range of z (the average number of neighbours of each individual). Due to the correlated local structure, a high-status individual now has more high-status individuals in its neighbourhood and consequently it is more difficult for the innovator to convince its neighbours to adopt, as according to equation 9.2 its relative influence is smaller. However, looking at the cascade size shown in Figure 9.6b we observe a similar dynamic to that seen in the uncorrelated situation. If the innovation spreads through 5 per cent of the population in a sufficiently dense network (average degree larger than four) then it spreads through almost the whole population.

Now we consider the situation where the innovation is introduced by a low-status individual. Figure 9.7a shows the spreading probability for the




Figure 9.7. Spreading probabilities of the high-status (light grey squares) and low-status subpopulations (dark grey dots) when the innovation is introduced by a low-status individual. We assumed $N=10000$ and $\phi_i=\phi=0.16$. (b) Corresponding average cascade size of the high-status (light grey squares) and low-status subpopulations (dark grey dots). (c) and (d) Same analyses as in (a) and (b) but with $N=1000$.

subpopulations of high-status (light grey squares) and low-status individuals (dark grey dots) for varying degree z of the network. At odds with the situation of no homophily where no or at most little spreading was observed (cf. light grey triangles in Figure 9.3a), we see that the innovation can spread in both the high- and low-status subpopulations. Again due to the correlated local structure, it is now easier for a low-status innovator to convince its neighbours to adopt the innovation. The neighbourhood of the innovator consists predominantly of other low-status individuals and therefore the innovator's relative influence on its neighbours is larger compared to the situation without homophily (see equation 9.2). Once the innovation has been adopted by a large number of low-status individuals, it can then spread even within the high-status subpopulation.

This temporal dynamic is illustrated in Figure 9.8a which shows the fraction of adopters in the high- (dashed light grey line) and low-status subpopulations (solid dark grey line) over time. Firstly we observe the S-shape of both adoption curves and secondly the described time-delay in adoption: the




Figure 9.8. (a) Fractions of the high-status (dashed light grey line) and low-status subpopulations (solid dark grey line) which have adopted the innovation at different time steps. (b) Corresponding fractions of whole population which have adopted the innovation at different time steps.




Figure 9.9. Spreading probabilities for the high-status (light grey squares) and low-status subpopulations (dark grey dots) for varying values of J , the measure of network correlation for (a) $N=10000$ and (b) $N=1000$.

innovation only starts spreading through the high-status population after almost all low-status individuals have adopted the innovation. Even though the spreading probabilities differ for both subpopulations we still see that if the innovation spreads through more than 5 per cent of the subpopulation the final sizes of the cascades are identical (cf. Fig. 9.7b). Interestingly, this dynamic causes a characteristic pattern of the adoption curve of the whole population. Figure 9.8b shows that the curve reaches a plateau before the high-status subpopulation starts adopting and then shows an increasing rate of transmission again.

Thus far we have kept the network correlation constant. However, in the following, we explore the relationship between the strength of homophily and the adoption dynamic. Figure 9.9 shows the spreading probabilities in the high-status (light grey squares) and low-status subpopulations (dark grey dots) for a network with average degree 5 (meaning that an individual has

on average five neighbours) and varying strengths of network correlation and therefore homophily for the situation when the innovation is introduced by a low-status individual.

Firstly, we observe that there exists a critical value of the parameter J (detailing the network correlation) which must be exceeded if the innovation is to spread if it is introduced by a low-status individual. An intuitive explanation of this behaviour is the following. If we rewrite the adoption decision (Eq. 9.2) of an individual i as:

$$\sum_{j \in \nu(i)} s_j \sigma_j < \phi \sum_{j \in \nu(i)} s_j \cong \phi z \bar{s}_i \quad (\text{Eq. 9.4})$$

where the sum $\sum_{j \in \nu(i)} s_j$ is approximated by the average degree z and the average status \bar{s}_i of i 's neighbours. A low-status neighbour alone is then able to convince individual i if its status is larger than $\phi z \bar{s}_i$. Now if individual i possesses a low status then an increase in J results in more low-status neighbours and consequently the average status \bar{s}_i decreases. Therefore, above a certain value for J the status of the innovator exceeds $\phi z \bar{s}_i$ and consequently individual i can adopt the innovation.

Secondly, strong homophilistic tendencies in the population cause different adoption dynamics in both subpopulations. While the innovation can spread in the low-status subpopulation relatively easily, it is far less likely to spread in the high-status subpopulation.

Another interesting aspect of the adoption dynamic is illustrated by the fact that the spread dynamic is highly dependent on population size. Figures 9.6c, 9.6d, 9.7c, 9.7d, and 9.9b show the results of the analyses described above for a network of $N=1000$ individuals. It is obvious that the effect of homophily is amplified in smaller populations. In particular, Figure 9.7c shows that in small populations it is unlikely that the high-status subpopulation will adopt an innovation if it is introduced by a low-status individual.

DISCUSSION AND CONCLUSION

In this chapter, we developed a simulation framework which relates individual-level processes, such as the decision to adopt an innovation or homophilistic bias, to population-level patterns and therefore enables us to study how changes in those underlying processes are reflected in, for example, the adoption curve. The analysis of the framework results in a set of hypotheses which can be tested against suitable archaeological data.

We considered a population where individuals are organized in a random Erdős-Rényi network and possess different levels of social status. The

innovation is introduced into the population by a single individual and the decision of other individuals to adopt depends solely on the number of individuals who have already adopted the innovation in the corresponding local neighbourhood, and their social status. We assumed explicitly that status differences are reflected by an individual's potential to influence other individuals. The higher the social status, the higher the influence of an individual. We have shown that heterogeneity in the population (in our case heterogeneity in social status) can alter the adoption dynamic of an innovation and pointed to the importance of the social status of the first adopter. We identified situations where innovations introduced by a high-status individual are likely to spread through the whole population, while innovations introduced by a low-status individual lead to almost no adoption at all. Interestingly, there is no difference in the adoption dynamic in the low- and high-status subpopulations; neither subpopulation is distinguishable on the basis of their adoption behaviour. Naturally, episodes where no spread occurs are not preserved in the archaeological record. However, based on the considered model we are able to quantify their probability.

Further, empirical research has shown that human populations often express a homophilistic bias, the tendency of individuals with similar traits to interact with each other more frequently than with individuals with dissimilar traits. In order to account for this behaviour, we relaxed the assumption that links between two individuals in the underlying network are drawn at random, independent of their status. Based on the social status differences, we assumed that individuals are more likely to interact with individuals with the same status and the resulting network is characterized by a correlated local structure. Varying the level of network correlation allowed us to investigate the consequences of homophilistic tendencies of different strengths on the adoption dynamic.

We found that if the innovation is introduced by a high-status individual, the population showed a similar adoption behaviour to that seen in the situation without homophily. However, if the innovation is introduced by a low-status individual we observed that the innovation can spread if the homophilistic bias is sufficiently high, and in this case that the spread dynamics differ in both subpopulations. While the innovation spreads in the low-status population with a high probability, it will only do so in the high-status population with a greatly reduced probability. Interestingly, we observed a time-delay in the adoption dynamic. The high-status subpopulation only starts adopting if a large fraction of the low-status subpopulation has already done so. This is reflected in the shape of the adoption curve of the whole population: it reaches a plateau before the high-status population starts adopting and then shows an increased rate of adoption again. We conclude that one possible explanation for such a population-level pattern is the differential adoption dynamic of different subpopulations. Further, we found

that the size of the population greatly influences the adoption dynamic. The smaller the population the stronger the effect of the homophilistic bias. In small populations it is almost impossible for the high-status subpopulation to adopt the innovation if it is introduced by a low-status individual, whereas the innovation is likely to spread in the low-status subpopulation.

In summary, our theoretical investigation has shown that the interplay between network topology and heterogeneity in the population can greatly influence the adoption dynamic of an innovation and therefore needs to be accounted for when explaining diffusion phenomena. Furthermore, the existence of a strong homophilistic bias in a population can result in different adoption dynamics in the low- and high-status subpopulations and therefore in a characteristic adoption curve of the whole population. Homophily can create and maintain cultural differences in a single heterogeneous population and the different adoption behaviours can act as a marker for the high- and low-status subpopulation.

We note that our analysis is based on Erdős-Rényi networks, which have been proven unrepresentative of human interaction networks. In recent years much effort has been devoted to understanding the empirical characterizations of real social networks (see, e.g. Albert and Barabási 2002), and subsequently different types of networks (e.g. small-world networks, scale-free networks) have been proposed to capture those patterns. For instance, it has been shown that scale-free networks reflect the observed power law degree distribution in networks such as the internet (Albert and Barabási 2002) or the network of movie actor collaborations (Barabási and Albert 1999). (Networks exhibiting a power law degree distribution are characterized by nodes whose numbers of neighbours are distributed according to a power law, i.e. the probability that a node has k neighbours is proportional to $k^{-\alpha}$, with α typically between 2 and 3. These networks are in particular characterized by the presence of few highly connected nodes.) Our future research is aimed at exploring the adoption dynamic on scale-free network topologies, which possess degree distributions that approximate more accurately the empirically observed degree distributions of social networks.

Additionally, recent advances in quantifying social networks using online-experiments or social media data (e.g. Centola 2010), as well as historical social networks (e.g. Mills et al. 2013; see in this volume Peeples et al. 2016; Evans 2016) will allow us to base the modelling framework on real-world data. In this way, the model can be used to infer theoretical expectations of, for example, the shape of the adoption curve based on (1) the assumed individual adoption dynamic (as quantified by equation 9.2) and (2) the social network. These theoretical expectations can then be compared to observed data using suitable statistical methods. We stress that especially with sparse archaeological data we do not expect the existence of a unique relationship between underlying

adoption decisions and observed population-level patterns; to the contrary, we expect that different mechanisms will be consistent with the observed data. Nonetheless, we anticipate that using a simulation framework such as the one developed in this chapter will help in narrowing down the range of possible adoption decisions that could have produced observed patterns, and thus will still be instructive in the face of uncertainty.

ACKNOWLEDGEMENTS

We thank Tom Brughmans, Anna Collar, and Fiona Coward for organizing the ‘The Connected Past’ conference in March 2012 in Southampton, UK which partly initiated this research. Further, we thank four anonymous reviewers for their comments which helped improving this paper. This research was supported by an Omidyar fellowship from the Santa Fe Institute to AK.

REFERENCES

- Albert, R., and Barabási, A. L. (2002). ‘Statistical Mechanics of Complex Networks’, *Reviews of Modern Physics*, 74:47.
- Barabási, A. L. and Albert, R. (1999). ‘Emergence of Scaling in Random Networks’, *Science*, 286:509.
- Bass, F. M. (1969). ‘A New Product Growth Model for Consumer Durables’, *Management Science*, 15:215–27.
- Bettinger, R. L., and Eerkens, J. (1999). ‘Point Typologies, Cultural Transmission, and the Spread of Bow-and-Arrow Technology in the Prehistoric Great Basin’, *American Antiquity*, 64:231–42.
- Burt, R. S. (1980). ‘Innovation as a Structural Interest: Rethinking the Impact of Network Position on Innovation Adoption’, *Social Networks*, 2:327–55.
- Burt, R. S. (1987). ‘Social contagion and innovation: Cohesion versus Structural Equivalence’, *American Journal of Sociology*, 92:1287–335.
- Centola, D. (2010). ‘The Spread of Behavior in an Online Social Network Experiment’, *Science*, 329:1194–7.
- Centola, D. (2011). ‘An Experimental Study of Homophily in the Adoption of Health Behavior’, *Science* 334:1269–72.
- Centola, D., Gonzales, J. C., Eguiluz, V. M., and San Miguel, M. (2007). ‘Homophily, Cultural Drift and the Co-Evolution of Cultural Groups’, *Journal of Conflict Resolution*, 51:905–29.
- Coleman, J. S., Katz, E., and Menzel, H. (1966). *Medical Innovation: A Diffusion Study*. New York, NY: Bobbs-Merrill Company.
- Dethlefsen, E., and Deetz, J. (1966). ‘Death’s Heads, Cherubs, and Willow Trees: Experimental Archaeology in Colonial Cemeteries’, *American Antiquity*, 31:502–10.

- Düring, M. (2016). 'How Reliable are Centrality Measures for Data Collected from Fragmentary and Heterogeneous Historical Sources? A Case Study', in Brughmans T., Collar A., and Coward F. (eds), *The Connected Past: Challenges to Network Studies in Archaeology and History*. Oxford: Oxford University Press, 85–101.
- Erdős, P., and Rényi, A. (1959). 'On Random Graphs', *Publicationes Mathematicae Debrecen*, 6:290–7.
- Evans, T. (2016). 'Which Network Model Should I Use? Towards a Quantitative Comparison of Spatial Network Models in Archaeology', in T. Brughmans, A. Collar, and F. Coward (eds), *The Connected Past: Challenges to Network Studies in Archaeology and History*. Oxford: Oxford University Press, 149–73.
- Feld, S. L. (1981). 'The Focused Organization of Social Ties', *American Journal of Sociology*, 86:1015–35.
- Granovetter, M. (1973). 'The Strength of Weak Ties', *American Journal of Sociology*, 78:1360–80.
- Granovetter, M. (1978). 'Threshold Models of Collective Behavior', *American Journal of Sociology*, 83:1420–43.
- Granovetter, M. (1982). 'The Strength of Weak Ties: a Network Theory Revisited', in: P. Marsden, and L. Nan (eds), *Social Structure and Network Analysis*. Beverly Hills: Sage, 201–33.
- Griliches, Z. (1957). 'Hybrid corn: An Exploration in the Economics of Technological Change', *Econometrica*, 25:501–22.
- Hägerstrand, T. (1967). *Innovation Diffusion as a Spatial Process*. Chicago: University of Chicago Press.
- Henrich, J. (2001). 'Cultural Transmission and the Diffusion of Innovations: Adoption Dynamics Indicate That Biased Cultural Transmission Is the Predominant Force in Behavioral Change', *American Anthropologist*, 104:992–1013.
- Kandler, A., and Steele, J. (2009). 'Social Learning, Economic Inequality and Innovation Diffusion', in M. J. O'Brien and S. J. Shennan (eds) *Innovation in Cultural Systems: Contributions from Evolutionary Anthropology Vienna Series in Theoretical Biology*. Cambridge, MA.: MIT Press.
- Katz, E., and Lazarsfeld, P. F. (1955). *Personal Influence: The Part Played by People in the Flow of Mass Communications*. Glencoe, IL: Free Press.
- Latané, B., and L'Herrou, T. (1996). 'Spatial Clustering in the Conformity Game: Dynamic Social Impact in Electronic Groups', *Journal of Personality and Social Psychology*, 70:1218–30.
- Lazarsfeld, P. F., Berelson, B., and Gaudet, H. (1968). *The People's Choice: How the Voter Makes Up His Mind in a Presidential Campaign*. New York: Columbia University Press.
- Lazarsfeld, P. F., and Merton, R. (1954). 'Friendship as Social Process: A Substantive and Methodological Analysis', in: M. Berger, T. Abel and C. Page (eds) *Freedom and Control in Modern Society*. New York: Van Nostrand.
- Lopez-Pintado, D., and Watts, D. J. (2007). 'Social Influence, Binary Decisions, and Collective Dynamics', Working paper, Institute for Social and Economic Research and Policy, Columbia University.
- Magee, J. C., and Galinsky, A. D. (2008). 'Social Hierarchy: The Self-Reinforcing Nature of Power and Status', *Academy of Management Annals*, 2:351–98.

- McPherson, J. M., and Smith-Lovin, L. (1987). 'Homophily in Voluntary Organizations: Status Distance and the Composition of Face-to-Face Groups', *American Sociological Review*, 52:370–9.
- McPherson, J. M., Smith-Lovin, L., and Cook, J. (2001). 'Birds of a Feather: Homophily in Social Networks', *Annual Review of Sociology*, 27:415–44.
- Mills, B. J., Clark, J. J., Peebles, M. A., Hass, Jr, W. R., Roberts, Jr, J. M., Hill, J. B., Huntley, D. L., Borck, L., Breiger, R. L., Clauset, A., and Shackley, M. S. (2013). 'Transformation of Social Networks in the Late Pre-Hispanic US Southwest', *Proceedings of the National Academy of Science*, 110(15):5785–90.
- Morrill, R. (1985). 'The Diffusion of the Use of Tractors Again', *Geographical Analysis*, 17: 88–94.
- Nadel, F. S. (1957). *Theory of Social Structure*. Glencoe, IL: Free Press.
- Noh, J. D. (2007). 'Percolation Transition in Networks with Degree-Degree Correlation', *Physical Review E*, 76:026116.
- Newman M. E. J., Watts D. J. and Stogatz S. H. (2002). 'Random Graph Models of Social Networks', *Proceedings of the National Academy of Science*, 99:2566–72.
- Peebles, M. A., Mills, B. J., Haas, W. R. J., Clark, J. J., and Roberts, J. M. J. (2016). 'Analytical Challenges for the Application of Social Network Analysis in Archaeology', in T. Brughmans, A. Collar, and F. Coward (eds), *The Connected Past: Challenges to Network Studies in Archaeology and History*. Oxford: Oxford University Press, 59–84.
- Richerson, P. J., Borgerhoff-Mulder, M., and Vila, B. J. (1996). *Principles of Human Ecology*. Boston: Pearson Custom Publishing.
- Rogers, E. M. (2003). *Diffusion of Innovations*, 5th edition. New York: Free Press.
- Rogers, E. M., and Shoemaker, F. F. (1971). *Communication of Innovations*. New York:Free Press.
- Ryan, B., and Gross, N. C. (1943). 'The Diffusion of Hybrid Seed Corn in Two Iowa Communities', *Rural Sociology*, 8:15–24.
- Scholnick, J. B. (2012). 'The Spatial and Temporal Diffusion of Stylish Innovation in Material Culture', *Advances in Complex Systems*, 15:1150010.
- Shalizi, C. R., and Thomas, A. C. (2011). 'Homophily and Contagion are Generically Confounded in Observational Social Network Studies', *Sociological Methods & Research*, 40:211.
- Shennan, S. J. (2011). 'Descent with Modification and the Archaeological Record', *Philosophical Transactions of the Royal Society B*, 366:1070–9.
- Steele, J. (2009). 'Innovation Diffusion and Travelling Waves', in S. J. Shennan (ed.) *Pattern and Process in Cultural Evolution*. Berkeley: University of California Press.
- Valente, T. W. (1993). 'Diffusion of innovations and policy decision-making', *Journal of Communication*, 43:30–45.
- Valente, T. W. (1995). *Network Models of the Diffusion of Innovations*. Cresskill, NJ: Hampton Press.
- Valente, T. W. (1996). 'Social network thresholds in the diffusion of innovations', *Social Networks*, 19:60–89.
- Van den Bulte, C., and Stremersch, S. (2004). 'Social Contagion and Income Heterogeneity in New Product Diffusion: A Meta-analytic Test', *Marketing Science*, 23: 530–44.

- Watts, D. J. (2002). 'A Simple Model of Global Cascades on Random Networks', *Proceedings of the National Academy of Science*, 99:5766–71.
- Watts, D. J., and Dodds, P. S. (2007). 'Influentials, Networks, and Public Opinion Formation', *Journal of Consumer Research*, 34:441–58.
- Young, H. P. (2009). 'Innovation Diffusion in Heterogeneous Populations: Contagion, Social Influence, and Social Learning', *American Economic Review*, 99:1899–924.

Index

- Actor-Network Theory 13, 24, 35
adjacency matrix 63, 154
adoption curve; dynamic, *see* diffusion:
 curve: dynamic
Aegean 16, 29, 129, 131, 135, 136–43,
 165–6
agency 15, 125, 126–9, 144
ariadne model 128, 135–41
- black box 22, 29, 128–9, 135
boundaries:
 categorical and conceptual 8, 45–7
 network 14, 59, 63, 70–4, 88
brokerage 89, 97, 99
Bronze Age 5, 16, 24, 28–9, 124, 134,
 136–41, 165
- centrality 5–6, 14–15, 17, 23, 27, 50, 63, 66–7,
 69, 73, 75–8, 85, 90–100, 177
 betweenness 92–6, 130, 137, 139, 160–1
 closeness 92–6
 degree 63, 67, 75–7, 92–6, 99
 eigenvector 63, 67, 69, 71, 73, 75–7, 92–6
 influence betweenness 160–1, 165–6
 PageRank 92–6, 160–1, 165–6, 170
ceramic, ceramics, pottery 21, 42–5, 50,
 60–2, 65–75, 78; *see also* *sigillata*;
 Tanque Verde red-on-brown
chaîne opératoire 30, 40–3, 51
clique 23, 117–19
clustering 90, 94, 117, 164
 hierarchical 166–7, 170
 spatial 16, 71, 77, 137
- coarse:
 grained 36
 -graining 29, 125, 129, 131–2, 142–3
 resolution 65, 67
 wares 43; *see also* ceramic
complete data, *see* data completeness
component 51, 129, 133
 network 69, 72–3, 106–7, 163
principal component analysis, *see* principal
 component analysis (PCA)
Crete 4–6, 24, 28–9, 136–9
cultural 22, 175, 177, 185, 193
 diffusion, transmission, *see* diffusion,
 cultural; transmission
 evolution 21, 175
- data completeness 14–16, 23, 28–30,
 59, 63, 104, 106, 108, 115,
 125, 130
degree 26, 70, 74, 75, 77, 91–9, 105, 108,
 117–19, 180–7, 190
 average degree 177, 182, 184, 188,
 190–1
 centrality, *see* centrality, degree
 degree distribution 177, 183, 186, 193;
 see also power law
 in-degree 92–6
 out-degree 92–6
density 177, 185
deterrence function 127, 131, 134–5, 141–3,
 150–1, 153–9, 168
- diffusion 4, 27, 193
 cultural; transmission 6, 16
 curve; dynamic; temporal 175–94
 innovation 175–7
- edge 9, 51, 62, 67, 78–9, 92, 126, 153–4,
 159–61, 163, 168–9
 dynamics 23
model 153–5, 159–60, 167–9
tie 9, 11, 50, 62–7, 71, 73–4, 78, 85–6,
 88–92, 94, 96–7, 99, 150, 177,
 181, 186
 social 37, 90, 92
 weak 4, 177
 weighted 62–3, 65–6, 69, 105, 109,
 111–12, 154–6, 160, 168–70
 unweighted 105, 154–5
- embeddedness 3, 23, 40, 43–4, 49, 88, 97–8
 spring-embedding, *see* spring-embedder
‘entanglement’ 24, 28, 30, 38, 47, 49
entropy 127, 133–4, 159
 maximization; maximum 29, 127, 159, 170
Erdős-Rényi networks, *see* network,
 random
- Gaul 41–4
Getty Museum 107–10
GIS 23, 60, 152
graph 63, 75, 79
graph theory 111, 123
gravity model, simple gravity model (SGM)
 29, 127–8, 133–5, 138–9, 141,
 155, 168

- Holocaust 86, 100
 heterogeneity:
 individual 176, 181, 183, 187, 192–3
 sociomaterial networks, *see* network:
 heterogeneity
 homophily 16, 135, 175–6, 181, 185–7, 189–93
 illicit antiquities trade 15, 103–8, 113, 115–16
 influence:
 influence betweenness, *see* centrality,
 influence betweenness
 social 67, 85, 90, 92, 94, 96–100, 177, 179,
 181–2, 186–7, 189, 192
 spatial 136, 141, 153
 stylistic 39, 42–3, 50
 innovation 5, 7, 16, 129, 175–94
 internet 38, 193
 Iron Age 41, 43, 124, 129, 134, 141, 143
 Jewish, Jews 14, 85–7, 92, 97, 100
 La Tène 42, 43
 Monte Carlo method 132, 186
 Neolithic 5, 26
 network:
 ‘beneficial’ 15, 126–7
 bipartite, two-mode 24–5, 64
 complete 15, 74
 directed 67, 79, 92, 105–13, 133, 135, 139,
 154, 155–6, 160, 168
 ego-network 88
 Erdős-Rényi, *see* network, random
 heterogeneity 23, 24, 26, 30
 ‘most likely’ 15, 126–7, 137
 one-mode 25, 64
 random 85, 177–8, 181, 182, 186, 191, 193
 scale-free 193
 ‘Small-world’ 23, 193
 social 3, 23–4, 27, 60, 63–4, 71, 79, 88, 90,
 99, 104, 115–16, 133, 135, 177, 183,
 185, 193
 sociomaterial, *see* network: heterogeneity
 support 14–15, 85, 89–90, 97, 99–100
 undirected 62, 105, 106, 107, 115–18, 133,
 139, 155, 160
 weighted, *see* edge: tie: weighted
 unweighted, *see* edge: tie: unweighted
 PageRank *see* centrality: PageRank
 Palaeolithic 25
 power law 155, 158, 180, 193
 principal component analysis (PCA) 164, 170
 proximal point analysis (PPA) 123, 133,
 153, 155
 random walk 160–1, 168–70
 reciprocity 22, 177
 Roman 13, 27, 35, 39, 41
 Second World War 14, 85–6
 shortest path 29, 96, 110, 112,
 114, 160–1
sigillata 13, 27, 35, 39–50
 similarity:
 Brainerd-Robinson similarity
 coefficient 62, 66
 matrix (matrices) 62, 69, 170
 measure 161–3
 metrics 66
 scores 62–3
 value 162, 166
 simple gravity model (SGM), *see* gravity
 model
 simulation 176, 180, 184, 191, 194
 social:
 influence, *see* influence, social
 media 193
 role 185
 status 16, 176, 181–2, 184–5, 191–2
 topology 35, 38, 39, 41, 44, 46–50
 social network analysis (SNA) 3, 11,
 15, 22–4, 27, 29, 36, 71,
 86, 112
 social physics 22, 24, 127
 spring-embedder 97
 Tanque Verde red-on-brown 69, 70, 79
 technology 30, 40, 43–4, 60, 65, 123,
 130, 133, 136, 138, 141–3, 165,
 167, 175
 Thera 132, 136–7, 139–41
 threshold 94, 96, 138, 153–5, 158–9, 164,
 177–9, 181–2
 tie, *see* edge
 topology:
 network 71, 73, 176, 193
 social 35, 38, 39, 41, 44, 46–50
 US Southwest 14, 59–60, 63–6, 69
 vertex (vertices) 150–1, 155, 160, 165,
 168–70
 visualize, visualization 11, 22, 62, 64,
 89, 94, 97, 100, 164, 166,
 169–70, 178
 weak tie, *see* edge: tie: weak
 weighting 65
 World-systems 25
 XTENT model 123, 135, 153, 170